Income in the Off-Season: Household Adaptation to Yearly Work Interruptions*

John Coglianese
Federal Reserve Board

Brendan M. Price
Federal Reserve Board

June 24, 2025

Abstract

Joblessness is highly seasonal. To analyze how households adapt to seasonal joblessness, we introduce a measure of seasonal work interruptions premised on the idea that a seasonal worker will tend to exit employment around the same time each year. We show that an excess share of prime-age US workers experience recurrent separations spaced exactly 12 months apart. Examining workers most prone to seasonal work interruptions, we find that they incur large earnings losses during the off-season that are little offset by other sources of income. On net, household income falls by about \$0.80 for each \$1.00 lost in own earnings.

Keywords: seasonality, seasonal employment, job loss, household income, household labor dynamics, unemployment, unemployment insurance

JEL codes: D10, E32, J63

*Email: john.m.coglianese@frb.gov and brendan.m.price@frb.gov. Price gratefully acknowledges financial support from the W.E. Upjohn Institute's Early Career Research Award program. We thank Marianne Bitler, David Card, Jeffrey Miron, Kathleen Mullen, Marianne Page, Giovanni Peri, Alexandra Roulet, Hannes Schwandt, Ashish Shenoy, Monica Singhal, Jenna Stearns, seminar participants at Bocconi University, Brigham Young University, INSEAD, UC Berkeley, UC Davis, UC Irvine, UC Merced, the US Census Bureau, and the Federal Reserve Board of Governors, and conference participants at the Society of Labor Economists, the Washington Area Labor Economics Symposium, and the All California Labor Economics Conference for helpful feedback. For information regarding the data and/or computer programs used for this study, please contact the authors at john.m.coglianese@frb.gov and brendan.m.price@frb.gov. The views expressed in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff.

Disclosure Statement: John Coglianese has no funding organizations, support, positions, or parties with the right to review to disclose. Brendan Price received financial support from a \$5,000 Early Career Research Award from the W.E. Upjohn Institute for Employment Research. Brendan Price has no other funding organizations, and no support, positions, or parties with the right to review to disclose.

Labor markets go through predictable seasonal downturns. Even in advanced economies, a slew of seasonal forces—such as winter weather, holiday shopping, and school recesses—induce systematic fluctuations at particular points in the calendar year. Figure 1 offers one illustration of this fact: the unemployment rate among prime-age US workers jumps by 0.7 percentage points in the typical January, when the post-Christmas drop in retail employment coincides with slack demand in construction, tourism, and other sectors reliant on favorable weather. The seasonal pattern of job flows is starker still: on average, the hazard rate of separating from a job into unemployment rises by 27% from December to January.[1] While economists may abstract from such "winter recessions" through seasonal adjustment, many workers are not so lucky. As one commentator has observed, "it does not help workers seeking jobs to tell them that seasonally adjusted they are employed" (Fromm, 1979).

This paper examines how households adapt to predictable seasonal work interruptions. To analyze how *aggregate* seasonality induces volatility in *individual* earnings, we introduce a measure of seasonal work interruptions premised on the idea that a seasonal worker is likely to exit from employment around the same time each year. Using three decades of panel data on US workers, we show that an excessive share of workers undergo repeated transitions from employment to non-employment spaced exactly 12 months apart. These recurrent transitions appear to be seasonal in origin, as they align closely with temporal and sectoral features of the seasonal cycle. To examine whether households recoup the income lost due to seasonal work interruptions, we use pre-separation worker characteristics to pinpoint which separators are most likely to separate again precisely one year later. On average, households prone to seasonal work interruptions replace only a small portion of their lost earnings: while seasonal drops in personal earnings are dampened by unemployment insurance, they are amplified by concurrent declines in partner earnings, so that the large majority of lost earnings passes through to lower household income.

[1] Own calculation using the Current Population Survey. Figure 1 plots unadjusted separation rates, but the change in flow rates noted in the text adjusts for cross-month differences in the spacing of successive reference weeks.

Our paper contributes to three main strands of existing literature. The first is a body of work that characterizes individual earnings dynamics in the wake of different kinds of job separations. Numerous papers have documented the slow process of earnings recovery after permanent job losses, such as those associated with plant closure (e.g., Jacobson, Lalonde, and Sullivan, 1993; Davis and von Wachter, 2011; Lachowska, Mas, and Woodbury, 2020). Other papers have catalogued the prevalence of temporary layoffs followed by workers' eventually being recalled to their former employers (e.g., Feldstein, 1975; Katz, 1986; Fujita and Moscarini, 2018; Nekoei and Weber, 2020). We situate seasonal work interruptions within the taxonomy of job loss, showing that seasonal separations are typified by a distinctive earnings process involving relatively rapid earnings recovery punctuated by recurrent drops in earnings as workers reenter their idiosyncratically timed off-seasons. Confirming in the US labor market a pattern previously noted in Austria (Del Bono and Weber, 2008; Nekoei and Weber, 2015), we find substantial overlap between seasonal jobs and ex post recalls, though neither phenomenon is a subset of the other.

Second, we contribute to the literature on the mechanisms through which households are partially insured against fluctuations in income. Prior work in this area has largely relied on annual data, which are unable to capture the seasonal volatility of income (e.g., Blundell, Pistaferri, and Preston, 2008; DeBacker et al., 2013). Recent work by Farrell and Greig (2015), Hannagan and Morduch (2015), and Morris et al. (2015) examines within-year volatility in household income but does not separate seasonal versus non-seasonal changes in income. Our focus on the margins along which households adapt to seasonal work interruptions contributes to the literatures on partial insurance from spousal labor supply (e.g., Lundberg, 1985; Stephens, 2002; Blundell, Pistaferri, and Saporta-Eksten, 2018) and from government transfers (e.g., Gruber, 1997; Dynarski and Gruber, 1997; Ganong and Noel, 2019). In contrast to prior research suggesting that an "added worker effect" mitigates the drop in household income incurred by job losers whose spouses are able to take on additional work, we find evidence of a *subtracted* worker effect in the context of seasonal separations.

Lastly, our paper relates to a literature on seasonal fluctuations in economic activity. Prior research has drawn connections between the seasonal and business cycles (Barsky and Miron, 1989; Geremew and Gourio, 2018), analyzed changes in seasonality over time (Sharpe and Smith, 2005; Gray and McDonald, 2010), and explored seasonal patterns in sectors like retail (Warner and Barsky, 1995) and housing (Ngai and Tenreyro, 2014). Like Price and Wasserman (2024), who document summer declines in women's employment and labor force participation, we examine the implications of seasonal work interruptions for individual workers and their households rather than for the aggregate economy. Though our focus is microeconomic, the window we open into seasonal fluctuations in household income may help inform our understanding of aggregate questions, such as the efficacy of stimulus policies carried out at different points in the calendar year (Olivei and Tenreyro, 2007).

The Scope for Household Adaptation

Seasonal workers face work interruptions around the same time each year. These interruptions may stem from seasonal fluctuations in labor demand, as in the case of construction workers laid off during the winter months. Alternatively, they may reflect labor supply factors, as in the case of parents who leave the labor force over the summer to care for young children (Price and Wasserman, 2024). In either case, seasonally recurrent work interruptions result in a reduction in earnings, forcing households to either replace the lost earnings, draw down savings, or reduce their consumption. Seasonal workers can (at least partially) insure themselves against lost earnings through several different channels.

First, workers could replace the lost earnings by taking up a different job. To the extent that a worker knows the date of separation in advance, they could search on-the-job and have a new job lined up by the time their seasonal job ends. On-the-job search tends to produce matches relatively quickly, as the typical worker searching while employed receives about three times as many offers per unit of search effort than the typical unemployed worker

(Faberman et al., 2022). However, on-the-job search may be less productive for seasonal workers, since their current industry is likely entering a down season and it may be difficult to switch industries. Since seasonal workers may be facing only a few months without their main job, they may find it less costly to simply remain non-employed during the off-season.

Second, a portion of the lost earnings may be replaced by transfer programs, particularly unemployment insurance (UI). Eligibility for UI is restricted to unemployed workers, as opposed to those who exit the labor force, and benefit amounts are typically based on earnings over a four-quarter base period. A priori, UI benefits could replace either a larger or smaller share of lost earnings for seasonal separators than for other separators. On the one hand, employers conducting seasonal layoffs may encourage their workers to apply for UI as part of an implicit contract, which could lead to high take-up rates among seasonal separators (Feldstein, 1978). On the other hand, not all seasonal separators remain in the labor force, and seasonal job losers may have patchier employment histories that limit their eligibility for UI; in addition, school employees—many of whom are laid off over the summer—are often statutorily ineligible for benefits. While we focus primarily on UI since it is the main transfer program through which laid-off workers replace lost income, seasonal separators may also be eligible to receive benefits from other programs, such as Supplemental Security Income (SSI), which insures disabled individuals with limited earnings histories; Temporary Assistance for Needy Families (TANF), which provides cash welfare payments to low-income families with children; or the Supplemental Nutrition Assistance Program (SNAP, or food stamps), which covers grocery expenditures for low-income households. Given up-front application costs and processing delays, however, these other programs may be poorly suited to seasonal workers who expect to be out of work for only a few months.

Third, when one member of a household experiences a seasonal separation, other members might increase their labor supply to replace the lost earnings. Indeed, prior research has documented an "added-worker effect" whereby secondary earners—typically women—enter the labor market or increase their hours when their partners lose work (Lundberg, 1985;

Stephens, 2002; Blundell, Pistaferri, and Saporta-Eksten, 2018). Since seasonal layoffs may be easier to predict than other layoffs, the added-worker channel could be more operative for seasonal layoffs than for the unanticipated layoffs typically examined by previous studies. However, since seasonal layoffs occur during a period of slack labor demand in the seasonal industry—and potentially in the broader labor market—other household members may find it particularly difficult to pick up additional work around the time of a seasonal layoff. Indeed, married partners are unusually likely to work in the same industry, and even for the same employer, and may be subject to the same seasonal downturn (Hyatt, 2019).

Lastly, seasonal workers could insure themselves against income loss by saving in advance of separation. While data limitations prevent us from studying savings or consumption among seasonal separators, prior work cautions that such anticipatory savings behavior may not be taking place. Stephens (2003, 2006) and Shapiro (2005) document that individuals do not perfectly insure themselves against the predictable income fluctuations that occur because paychecks and transfer payments are distributed on particular days of the month. Ganong and Noel (2019) and Gerard and Naritomi (2021) show that job losers accumulate relatively little excess savings in the lead-up to UI benefit exhaustion, another predictable event. If consumption largely tracks income, as these studies suggest, then our estimates of seasonal separators' ability to smooth household income are informative about whether they are able to insure themselves against the loss of earnings.

Data

To identify which households are exposed to seasonal work interruptions, and to analyze how they adapt to such interruptions, we employ two nationally representative household survey datasets that are available on a monthly basis over the years 1984–2013.[2] The Current Population Survey (CPS) furnishes large samples that are well-suited to identifying sea-

[2]While the CPS is available over a longer timespan, the SIPP was launched in 1984 and switched to annual interviews in 2014, limiting its usefulness for analyzing seasonal changes in income after that point.

sonal workers. The Survey of Income and Program Participation (SIPP) allows us to track households' employment, earnings, and non-labor income across multiple seasonal cycles. Appendix A offers additional detail about how we construct each sample.

CPS data

The CPS is a household survey conducted monthly by the Bureau of Labor Statistics, averaging about 105,000 adults per month. Households selected for the sample are interviewed for four consecutive months, dropped from the sample for the next eight months, and then interviewed for a final four consecutive months. Conveniently for our purposes, this 4–8–4 interview design enables us to match respondents both month-to-month and year-over-year, with respondents observed in the same portion of the calendar year on both occasions that they come into view. As explained in the next section, we proxy for seasonal work interruptions using recurrent separations into employment spaced 12 months apart.[3]

Using CPS extracts spanning 1984–2013, we construct a monthly panel of individuals ages 25–54, so as to focus on seasonal separations experienced by people in their prime working years. We exclude younger workers because periodic separations among young adults likely reflect seasonal fluctuations in schooling opportunities as well as in job availability. We exclude older workers so as to abstract from retirement decisions, which may be timed to coincide with summer vacations or with the seasonal cycle in the housing market.

The CPS reports each respondent's employment status during each month's reference week. We code each respondent as employed for a given month if the respondent either (i) was at work during the reference week or (ii) held a job but was not at work during the

[3]The CPS does not provide direct information for the full sample on whether workers are employed in seasonal jobs. Some information is available for limited subsamples: (i) non-participants who have worked in the last five years are asked why their last job ended, and one possible response is "seasonal job completed"; (ii) respondents to the Displaced Worker Supplement can report that their separation was due to a seasonal job ending; and (iii) respondents to the Contingent Work Supplement who report working temporary jobs are asked why their job is temporary, and one option is "nature of work/seasonal". Each of these subsamples is likely to omit short-lived work interruptions, and it is not clear that respondents' notions of "seasonal" jobs line up with the focus of this study. For these reasons, we view proxying for seasonal work interruptions with recurrent separations as more informative than using any of these subsamples.

reference week. We code a worker as experiencing a separation if the worker was employed in the previous month but not in the current month. The CPS separation concept excludes brief non-employment spells contained entirely between reference weeks, as well as job-to-job transitions that are unaccompanied by an intervening spell of non-employment.

Appendix Table F.1 reports summary statistics for our CPS sample, separately for employed workers, workers who have just separated into non-employment, and the subset of separators who experience a second separation exactly 12 months after the first one. Among separators, annual repetitions are disproportionately common in agriculture/fishing/forestry, construction, and educational services, three sectors that exhibit pronounced seasonal employment fluctuations. They are underrepresented in healthcare, a sector with comparatively stable employment levels throughout the calendar year.

While the CPS is well suited to identifying seasonal workers, it is uninformative about how earnings and income fluctuate within households from month to month, as each household reports these outcomes only once per year as part of the outgoing rotation group module. For this, we turn to the SIPP.

SIPP data

The SIPP is a household survey produced by the US Census Bureau, which interviews about 40,000–70,000 adults each month. It is structured as a series of longitudinal panels, with all respondents in a panel commencing their interviews in the same year and continuing for between 2.5 and 5 years (depending on the panel). We exploit the SIPP's long, unbroken panels to track employment, earnings, and income dynamics after separations from employment.[4]

The SIPP records labor force status on a weekly basis. Monetary receipts—including individual earnings, household income, and line items for a multitude of government transfer programs—are instead tallied on a monthly basis. To align our (weekly) employment data

[4]Like the CPS, the SIPP does not directly measure the extent of seasonal work in the full sample. Business owners whose business recently ended are asked why, and one possible response is that the "season ended for a seasonal business", but otherwise the SIPP collects no information about seasonal work.

with our (monthly) income data, we first code a job separation as occurring in some week w if an individual reported being employed in week $w-1$ but non-employed in week w.[5] We then aggregate these weekly codes to the monthly level by recording whether each individual experienced at least one separation during a given month.[6] In our SIPP-based analyses, our notion of separations thus encompasses all new jobless spells that last for at least one week.

We construct a monthly panel of prime-age SIPP respondents over 1984–2013, matching our CPS sample. Appendix Table F.2 reports summary statistics for our SIPP sample. As in the CPS, annually recurrent separators are overrepresented in highly seasonal industries. Relative to separators as a whole, annually recurrent separators have below-average earnings and household incomes in the month prior to separation, hinting that seasonal workers may be negatively selected on baseline earnings potential.[7]

Identifying Seasonal Work Interruptions

We identify seasonal work interruptions on the basis of annual periodicity in individual work histories. First, we show that workers who separate from employment in a given month are disproportionately likely to separate again exactly 12 months later, enabling us to isolate a component of employment seasonality that is predictable at the individual level. Second, we use a machine learning algorithm to pinpoint workers who are prone to such annually recurrent separations. Our proxy for seasonal work interruptions aligns closely with seasonal patterns in the labor market across sectors, calendar months, and climates.

[5] The SIPP assigns each individual to one of five labor force statuses. We code individuals as "employed" in a given week if they are either *at work* or *absent from work but not on layoff*; as "unemployed" if they are *absent on layoff* or *jobless and looking for work*; and as "non-participants" if they are *jobless and not looking for work*.

[6] For instance, a worker who was employed in the last week of December but unemployed in the first week of January will be coded as having a separation in January. If the worker first returns to work in mid-March, then is laid off and exits the labor force in mid-April, the worker has a separation in April but no separation in February or March.

[7] Prior work has found that seasonal workers receive positive compensating wage differentials for the disamenity of seasonal work interruptions (Moretti, 2000; Del Bono and Weber, 2008). If the recurrent job separators in our sample receive similar compensation, then their latent earnings potential may be lower than the observed earnings in Appendix Table F.2 imply.

Annually Recurrent Separations

Because different seasonal phenomena occur at different points in the year, seasonality in aggregate labor market indicators does not fully capture the seasonal shifts in work opportunities or life circumstances experienced by individual workers. For example, an agricultural laborer might work only during the summer months, while a school bus driver might work only during the school year. Even within industries, workers' exposure to seasonality will vary with occupation, geography, and idiosyncratic factors. For example, a construction worker in Minnesota may be laid off at the start of winter, while a similar worker in Arizona may be able to work year-round. To analyze household adaptation to yearly work interruptions, we must first identify which specific workers are at risk of seasonal separations.

Our method is predicated on the idea that seasonal workers will tend to experience recurrent separations around the same time each year, such as the start of winter, the end of the school year, or the end of the ski season. Concretely, we take a sample of workers who separate from employment at some point in time and calculate their likelihood of experiencing additional separations in subsequent months. Seasonal workers should be excessively likely to experience another separation exactly 12 months later, while non-seasonal separators will be no more likely to experience a separation at an annual horizon than at other, similar spans of time.[8] Therefore, the *excess recurrence* of separations at 12-month spans provides an estimate of seasonal work as experienced by individuals.[9]

For a worker i who separates from employment into non-employment in year-month t_0, let ρ_τ be the probability of experiencing another separation at date $t_0 + \tau$. By estimating ρ_τ over a range of $\tau \in \{1, \ldots, \tau_{\max}\}$, we track the evolution of this rate of recurrence following a baseline separation. To measure the excess recurrence of separations at an annual horizon, we

[8] To formalize the intuition underlying our approach, Appendix B describes a stylized economy in which a mixture of seasonal and non-seasonal workers cycle in and out of work.

[9] Precursors to our approach appear in de Raaf, Kapsalis, and Vincent (2003) and Del Bono and Weber (2008), who define seasonal jobs as those which see separations during the same three-month portion of the calendar year in multiple consecutive years. We depart from these earlier papers both in our focus on household adaptation to seasonal separations and in our approach of pinpointing which separators are *likely* to experience recurrent separations.

must estimate the counterfactual recurrence rate at this horizon in the absence of seasonality. A natural benchmark is to use the separation probabilities at the neighboring horizons $\tau = 11$ and $\tau = 13$, appealing to smoothness of the underlying non-seasonal dynamics. Subtracting the average of the 11th and 13th recurrence probabilities from the 12th yields a measure of seasonality as experienced by individuals:

$$\text{excess recurrence} \equiv \rho_{12} - \frac{1}{2}(\rho_{11} + \rho_{13}) \tag{1}$$

In the CPS, the 4–8–4 interview structure allows us to observe recurrent separations for $\tau \in \{10, 11, 12, 13, 14\}$, but not at other horizons. In the SIPP, we follow workers for a year and a half after the baseline separation, $\tau \in \{1, 2, \dots, 18\}$. As such, we can estimate excess recurrence in both datasets.[10]

Figure 2 plots the time path of recurrent separations in the wake of an initial separation from employment into non-employment.[11] In both datasets, the probability of a recurrent separation spikes 12 months after baseline, indicating that separations tend to recur at annual frequencies. We estimate an excess recurrence of 1.4 p.p. in the CPS and 2.0 p.p. in the SIPP. Appendix C shows that these results are robust to a suite of alternative specifications and sample restrictions.

In Appendix D, we show that our measure of excess recurrence aligns with three characteristic features of labor market seasonality. First, excess recurrence is largest among separations that occur at the start of winter or the start of summer, when the employment-to-population ratio among prime-age individuals typically declines (Appendix Figure F.2). Second, excess recurrence is especially pronounced in sectors like agriculture, educational

[10]To implement this procedure, we count the number of separations S_i experienced by each individual i, stack S_i copies of i's data, associate each copy with a base separation at t_0, and retain all available observations from $t_0 + 1$ through $t_0 + \tau_{\max}$. We then estimate $\{\rho_\tau\}$ in a single, stacked regression. We cluster standard errors at the individual level, both to allow for serial correlation in individuals' realized job separations and to account for overlapping periods in our stacked data. Our approach treats all observed separations symmetrically, rather than treating some as base separations and others as outcomes, and ensures that our results are representative of all job separations.

[11]Note that it is possible (and indeed fairly common) for SIPP respondents to register job separations in back-to-back months, since our measure of separations is based on underlying weekly detail.

services, and entertainment/recreation that exhibit pronounced seasonal fluctuations (Appendix Figure F.3). Third, workers who undergo a temporary layoff are especially likely to experience a repeat separation 12 months later (Appendix Figure F.4), often from the same employer (Appendix Figure F.5). The large role played by temporary layoffs is consistent with workers being furloughed during the off-season with the expectation of being brought back to work when seasonal conditions improve.

Pinpointing seasonal separators

Our estimates of excess recurrence capture the extent of seasonal separations *ex post*. We now identify workers who are *at risk* of experiencing a recurrent separation based on a set of characteristics observable at the time of their first separation. Our method is able to capture subtle differences in seasonality across jobs—recognizing, for instance, that construction jobs are seasonally interrupted only during winter months in cold states. Ranking separators by our measure of their risk of separating again, we find that those in the top decile have an average excess recurrence rate more than seven times higher than those in the bottom half.

We start by estimating the variation in recurrent separations across different types of jobs. For a worker in job type j who experiences an initial separation at time t_0 with job characteristics X_j, we denote the probability of experiencing a separation exactly τ months later as $\rho_\tau(j) \equiv \Pr\left(\mathrm{Sep}_{t_0+\tau} \mid \mathrm{Sep}_{t_0}, X_j\right)$. We treat this probability as an unknown (possibly non-linear) function f of both the job's characteristics and the observation horizon, that is, $\rho_\tau(j) = f(X_j, \tau)$. The excess recurrence rate for jobs of type j is equal to

$$\widetilde{\rho}(j) \equiv f(X_j, \tau = 12) - \frac{1}{2}\left(f\left(X_j, \tau = 11\right) + f\left(X_j, \tau = 13\right)\right). \tag{2}$$

One way to estimate f would be to divide values of X_j into discrete cells and estimate recurrent separation probabilities separately within each cell. For instance, if X_j includes industry and occupation, this would amount to computing excess recurrence separately within each unique combination of industry and occupation. However, while this approach is completely non-parametric, it may give very noisy estimates when the number of cells is large and consequently many cells contain only a few observations. A major limitation of this approach is that many cells will have true excess recurrence rates similar to those in "nearby" cells with similar job characteristics, but this information is ignored in estimation.

Instead of dividing into cells, we estimate the function f non-parametrically using gradient-boosted trees, a commonly used machine-learning algorithm for estimating high-dimensional functions with unknown functional form.[12] This algorithm builds up an estimate of the unknown function by iteratively applying a series of small decision trees, with each tree in the sequence being fit to the residuals from the previous trees and weighted according to the gradient of the objective function. This process of fitting many weak learners to the data, at each step attempting to correct the largest mistakes from previous steps, helps deliver accurate estimates and avoid overfitting, since it represents a form of gradient descent over a function space (Mason et al., 2000).

For characteristics X_j, we focus on the most important variables in the context of recurrent separations. To account for sectoral differences in the prevalence of seasonal work, we include the industry and occupation in which the individual worked in the month preceding the initial separation. To account for weather-related differences in the timing and amplitude of the seasonal cycle, we also include the historical average maximum daily temperature in January in the individual's state of residence (see Appendix A for details on the temperature data). Lastly, we include the number of months since the initial separation τ as well as the calendar month $(\text{January, February, } \dots)$ in which the original separation occurred.[13] We emphasize that gradient-boosted trees naturally allow for flexible interactions among the input variables, without the econometrician needing to pre-specify which interactions ought

[12]This algorithm was developed by Breiman (1997) and Friedman (2001, 2002). Friedman, Hastie, and Tibshirani (2001) provide a helpful overview of the method.

[13]Industry and occupation are treated as categorical variables, while temperature and τ are ordinal. To account for the "circularity" of the calendar cycle, we express each calendar month as a pair of trigonometric coordinates, which are then each treated ordinally. The restriction of ordinality is relatively minimal since gradient-boosted trees do not impose linearity, continuity, or monotonicity on ordinal covariates.

to be included in the model.

We estimate the function f using our CPS sample.[14] Before estimating with the full sample, we use cross-validation to select the key parameters of our algorithm, with the aim of maximizing the out-of-sample performance of our predictions.[15] Minimizing cross-validated mean squared error selects 800 trees with a learning rate of 0.005, which we then use as the parameters for estimating on the full CPS sample.

We then construct predicted excess recurrence for all separators in the SIPP using the estimated function $\hat{f}$. Specifically, for an individual i with baseline characteristics j, we construct the predicted excess recurrence rate $\hat{\rho}_i$ as:

$$\hat{\rho}_i = \hat{f}\left(X_j, \tau = 12\right) - \frac{1}{2}\left(\hat{f}\left(X_j, \tau = 11\right) + \hat{f}\left(X_j, \tau = 13\right)\right) \tag{3}$$

This forms an estimate of the extent to which we would predict an individual who is separating from employment today is excessively likely to separate again exactly 12 months in the future. By design, this measure is based only on information that is known at the time of the initial separation, so that it does not depend on subsequent outcomes of interest.

We use this measure to separate workers into those most at risk of seasonal separations (top decile) and least at risk (bottom half).[16] Table 1 reports several characteristics of these two groups of separators. While similar on a number of dimensions, the two groups have a notably different distribution across industries. The most seasonal separators are far more likely to work in agriculture or education, and far less likely to work in healthcare. However,

[14] By training our predictive model on the CPS, rather than the SIPP itself, we avail ourselves of the larger CPS sample for the data-hungry estimation step, and we sidestep the need to divide our SIPP data into estimation and prediction samples. This allows us to deploy our full SIPP sample to the analysis of household adaptation.

[15] For each possible choice of the parameters, we divide the sample into 5 equally sized subsamples. For each subsample, we estimate the function f on the other 4/5 of the data, construct predicted values for the subsample, and compute the mean squared error (MSE) of these predictions. We repeat this for each possible choice of the parameters on a grid with all combinations of the number of trees in the set $\{50, 100, 200, 400, 600, 800, 1000\}$ and learning rates in the set $\{0.05, 0.01, 0.005, 0.001\}$ and select the combination with the smallest MSE. All choices use trees with 63 leaves and at least 10 observations per leaf.

[16] Throughout the paper, we use sampling weights when calculating decile cutoffs, to ensure that each decile represents one-tenth of the population of job separations.

our algorithm is not solely relying on industry to categorize workers, as evidenced by both groups having similar shares of construction employment.

To illustrate the kind of subtle interactions that our algorithm is able to detect, Figure 3 shows the share of construction workers whom our algorithm assigns to the top decile of predicted excess recurrence, separately by month of separation and by residence in the coldest or warmest tercile of states. Consistent with our prior intuition that seasonal layoffs in construction are commonly associated with the onset of adverse winter weather, our algorithm classifies construction workers as "seasonal" only if they exit employment around the start of winter, and only if they do so in a cold location.

Our algorithm also picks up on interactions between industry and occupation, as shown in Appendix Table F.4. For example, within agriculture, separations from transportation occupations are nearly ten times more likely to be classified as seasonal than separations from administrative support. Construction and entertainment/recreation also show large differences across occupations. Within education, by contrast, professional occupations (such as teachers) and other occupations (such as administrative assistants) exhibit a similar likelihood of appearing in the top decile. Occupational heterogeneity also interacts with timing within the year: seasonal occupations in agriculture and construction are more likely to separate in the winter, while education workers are more likely to separate in the summer.

As an out-of-sample test of our predictive algorithm, we rank individual separators in the SIPP by their predicted excess recurrence and measure the actual excess recurrence within each decile of this ranking. If our predictions are valid, we should expect that the deciles with higher *predicted* excess recurrence also have higher *estimated* excess recurrence. Appendix Figure F.6 shows that this is indeed the case: the estimated average excess recurrence rate is 6.1% for those in the top decile, more than seven times larger than the average rate of 0.8% for those in the bottom half.

Earnings and Income Dynamics

Having established our method for identifying seasonal work interruptions, we now examine their consequences for earnings and household income. We use an event study approach to examine how earnings and income evolve for separators in our sample over the year and a half following their initial separation. We find that separators predicted to be the most seasonal do indeed exhibit income dynamics consistent with recurrent seasonal separations, whereas separators predicted to be least seasonal do not. Our central result in this section is that lost earnings from seasonal separations largely pass through to overall household income.

We focus on three outcomes of interest. First, we estimate the change in personal earnings, which represents the loss of income directly from the seasonal job separation. Second, we estimate the change in household income, which captures the response of income along all margins at the time of seasonal work interruptions. Lastly, we estimate the pass-through ratio formed by dividing the change in household income by the change in personal earnings. This ratio—which measures the fraction of the income lost from a seasonal job ending that is not made up by other sources of income within the household—summarizes the degree to which households adapt to seasonal work interruptions. We defer to the next section a look at the margins that account for this adaptation.

For a sample of N individuals who are in either the top decile or bottom half of predicted excess recurrence, with t_0 denoting the month of the initial separation, we measure the evolution of an outcome y observed τ months after the initial separation as

$$\beta_y^{(\tau)} \equiv \frac{1}{N} \sum_{i=1}^{N} \frac{y_{i,t_0+\tau} - y_{i,t_0+\tau-1}}{\text{Earnings}_{i,t_0-1}} \tag{4}$$

Thus $\beta_y^{(\tau)}$ is the average month-over-month change in y in a given group of separators, with changes denominated by the separator's earnings one month prior to the initial separation. Normalizing all outcomes by earnings allows us to compare different margins of income loss (or recovery) in the same units. To reduce the influence of outliers, we exclude individuals with pre-separation earnings less than \$450/month in January 2017 dollars, and we winsorize month-to-month changes in household income so that each is between -300% and 300% of the separator's pre-separation earnings.[17,18] We adjust changes in earnings and (later) other components of income proportionately to ensure that they sum to the winsorized total.

We add up our estimates $\beta_y^{(\cdot)}$ to get the cumulative change in y over the n months following the initial separation:

$$\widetilde{\beta}_y^{(n)} \equiv \beta_y^{(0)} + \beta_y^{(1)} + \ldots + \beta_y^{(n)} \tag{5}$$

We sum up month-to-month changes (rather than directly estimating long differences from t_0 to $t_0 + n$) because doing so allows us to retain each respondent in the sample for as many periods as possible in cases where they ultimately attrit from the sample. This affords us a larger sample for identifying each month-to-month change and mitigates any potential biases stemming from non-random attrition. Taking long differences yields similar results.

Note that the coefficients $\widetilde{\beta}_y^{(n)}$ do not represent the *percent change* in the outcome y: instead, they are the change in y *as a percentage of pre-separation earnings*, which enables more straightforward comparisons. For example, consider a household composed of a seasonal worker who earns \$1,000 per month and a non-seasonal worker who earns \$3,000 per month, with no other sources of income, giving a total household income of \$4,000. Suppose the seasonal worker is laid off and this income is not replaced by any other source, so that the worker's earnings fall by \$1,000. This is both a 100% decrease and a decrease as a percent of pre-separation earnings of 100 percentage points (p.p.). However, although household income has only fallen by 25%, the \$1,000 decrease in household income *as a percentage of the seasonal worker's pre-separation earnings* is 100 p.p., the same as for personal earnings. The fact that both variables fall by the same amount when expressed in terms of the same base indicates complete pass-through, which might otherwise be difficult to discern from a

[17] Our \$450 threshold for (deflated) pre-separation monthly earnings corresponds to 80 hours of minimum-wage work circa 1990, the point in our sample period when the real federal minimum wage reached its nadir.

[18] A total of 2.73% of observations fall outside of $[-300\%, 300\%]$ and require winsorizing.

comparison of the percent changes in these variables.

Applying this procedure, Figure 4 shows the estimated changes in total personal earnings for individuals who are in either the top decile or bottom half of predicted excess recurrence (left panel), along with changes in household income for each (right panel). Because both variables are normalized by pre-separation earnings, they can be compared on the same basis. If the change in the separator's personal earnings is not offset (or exacerbated) by any other components of household income, then the two graphs will show identical changes.

The estimates shown in Figure 4 demonstrate that lost earnings from seasonal jobs pass through to household income nearly one-for-one. Among the most seasonal (top-decile) separators, who are the most likely to separate from a seasonal job 12 months after the initial separation, personal earnings as a percent of pre-separation earnings decline by 18.6 p.p. between event-times $\tau = 10$ and $\tau = 13$ months; household income (in comparable units) declines 17.3 p.p. over the same period.[19] This amounts to a decline in household income of about \$0.93 for every \$1 of lost earnings from a seasonal job. The high rate of pass-through implies that only a small portion of lost earnings is recouped through increases in the other components of household income, including the earnings of other individuals within the household, government transfers, and other income from non-labor sources. The substantial drops in earnings and income for the most seasonal separators about a year after their initial separation stand in contrast to those of the least seasonal (bottom-half) separators, who experience a 2.7 p.p. increase in earnings and a 2.0 p.p. increase in household income over the same time period.

These comparisons relate to earnings and income *received in month* $\tau = 13$ relative to month $\tau = 10$. To get an estimate of households' *total* foregone income over the off-season, we sum the shortfall in each household's income receipts in months $\tau \in \{11, 12, 13\}$, relative to its receipts in month $\tau = 10$: that is, for each household i we calculate the quantity

[19] The 18.6 p.p. drop in earnings is notably larger than the 6.1% excess recurrence rate for the most seasonal separators shown in Appendix Figure F.6, indicating that some workers who do not experience a transition from work into non-work nonetheless experience a decline in earnings.

$\Delta Y_i \equiv \sum_{\tau=11}^{13} \frac{y_{i,t_0+\tau} - y_{i,t_0+10}}{\text{Earnings}_{i,t_0-1}}$. In Appendix Figure F.10, we present the empirical cumulative distribution function of ΔY_i for households in the top decile and bottom half of predicted excess recurrence. Among the most seasonal households, 28.1% face total off-season income losses exceeding one month of the separator's pre-separation earnings (compared to only 20.3% of the least seasonal households over the same period). Roughly one in six seasonal households face total losses exceeding two months of prior earnings; roughly one in nine lose income equivalent to three months' worth of earnings. These calculations highlight that an appreciable share of seasonally exposed households face off-season income losses on the order of multiple months' paychecks.

Margins of Adaptation

We now turn to the margins along which households recoup (or do not recoup) the earnings lost to seasonal work interruptions. We do so by charting the evolution of various components of household income between 10 and 13 months after an initial separation, separately for those most and least prone to see a recurrent separation 12 months after the initial one.

We start by examining the focal separator's earnings in more detail, decomposing the change in earnings between the original and other industries. We then consider changes in income earned by other household members, notably the separator's spouse or unmarried partner. We also evaluate the role of government transfers in replacing lost earnings from seasonal work interruptions. Lastly, we show how responses along these margins vary with observable characteristics of seasonal separators, so as to explore heterogeneity in household adaptation to seasonal earnings losses.

Separators' earnings and employment

We start by determining which kinds of jobs are proximately responsible for the earnings decline. We divide earnings into two categories based on the worker's industry prior to the baseline separation (i.e., the industry affiliation at event-time $\tau = -1$): earnings in the original industry and earnings in all other industries. For the 1990–1993 panels, which have reliable job identifiers, we additionally separate earnings in the original industry into earnings at the original employer and earnings at all other employers in the original industry. Table 2 shows how earnings at each of these kinds of jobs evolve from 10 to 13 months after an initial separation for the most and least seasonal separators.

Strikingly, we find that the overall earnings decline for the most seasonal separators is entirely accounted for by declines in earnings in the original industry. The 1990–1993 panels further reveal that this decline occurs almost entirely within the original employer, consistent with our earlier finding that many seasonal workers separate from the same job in back-to-back years. Earnings at other jobs in the original industry also contribute to the decline, albeit to a smaller extent.[20] Earnings growth in other industries, though positive, is small, statistically insignificant, and sluggish in relation to the growth rates registered by the least seasonal separators. Perhaps surprisingly, few seasonal separators appear to be changing industries to pursue "off-season" jobs pending the next seasonal upswing.

Partner and other household members' earnings

Next, we examine whether earnings by partners and other household members offset some of the lost earnings from seasonal separations. As previous work has documented, when a married individual separates from a job there is often a contemporaneous increase in their partner's labor supply (Lundberg, 1985; Stephens, 2002; Blundell, Pistaferri, and Saporta-Eksten, 2018). This "added worker effect" offsets some—though typically not all—of the household income lost due to the job separation.

In contrast to much of the previous literature, we find that partner earnings tend to *fall* during recurrent seasonal separations. Figure 5 shows that partner earnings for the

[20] As an example of a work history that would yield this pattern, a worker might be laid off from a construction firm in the winter in year 1, obtain a different construction job in the spring, and then separate from this new job in the winter of year 2.

most seasonal separators decline about a year after the initial separation, primarily in the separator's industry. Partner earnings as a percent of the focal separator's pre-separation earnings fall by 2.3 p.p. for earnings in the separator's industry and 0.3 p.p. in all other industries, compared with an increase of 0.1 p.p. and an increase of 0.9 p.p., respectively, for the least seasonal separators (see Appendix Table F.5). This pattern indicates a "subtracted worker effect" whereby changes in spousal labor supply exacerbate the income loss from a seasonal separation rather than providing insurance and partially offsetting the lost earnings. The fact that this subtracted worker effect is driven by earnings in the same industry as the seasonal separator's initial job suggests that partners are likely subject to the same seasonal fluctuations.[21] For seasonal workers, having a partner who works in the same industry creates correlated income volatility over the year, diminishing the scope for this margin to smooth out fluctuations in one's own income.

The difference between this result and the previous literature appears to reflect our focus on seasonal separators, rather than differences in our approach or methodology. As seen in Figure 5, while the most seasonally exposed separators experience a subtracted worker effect, the least seasonally exposed separators see the same added worker effect in the wake of their initial separation as has been documented previously. The added worker effect for the least seasonal separators is driven by earnings in other industries besides the separator's own, consistent with partner earnings serving as a form of partial insurance for these separations. This suggests that the correlation of income dynamics between partners makes seasonal separators distinct from other types of job losers by depriving them of an often-utilized form of intra-household insurance.

We also examine whether there are further changes in labor supply within households in response to seasonal separations. Appendix Table F.5 reports changes in the earnings of other, non-partner household adults for the most and least seasonal separators. Earnings of other household members change little for both groups 10 to 13 months after the base

[21] Hyatt (2019) shows that many couples who work in the same industry in fact work at the same establishment, potentially increasing the correlation of earnings between partners working in seasonal jobs.

separation, indicating that this margin does not much mitigate or exacerbate lost earnings from seasonal work interruptions.

Government transfers

A final margin through which households may replace lost earnings is the receipt of government transfers. We focus on three types of transfers: unemployment insurance (UI), means-tested cash transfers (SSI, TANF, etc.), and the Supplemental Nutrition Assistance Program (SNAP, a.k.a. "food stamps"), which together form the backbone of the US social safety net for displaced workers and low-income households. We trace how the amount of income or probability of receipt evolves in the wake of an initial separation for the most and least seasonal separators.

Our estimates indicate that UI provides a meaningful, though far from total, degree of income-replacement for seasonal separators. We find that total household UI receipts as a percent of pre-separation earnings rise by 4.7 p.p. among the most seasonal separators between 10 and 13 months after an initial separation, compared to no change for the least seasonal separators.[22] This offsets approximately one-quarter of lost earnings.

Although the SIPP is "specifically designed to determine eligibility and receipt of government transfers" (Meyer, Mok, and Sullivan, 2015), households may nonetheless underreport transfer receipts. Meyer, Mok, and Sullivan compare self-reported receipt of UI benefits (and other transfers) in the 1983–2012 SIPP to corresponding administrative aggregates, finding that as much as 30% of UI dollars may go unreported in the SIPP. We scale up the reported UI benefits received by each household in the SIPP using the year-by-year estimates of underreporting from Meyer, Mok, and Sullivan (2015) and recompute all of our estimates. Adjusted for underreporting, Figure 6 shows that UI receipts rise by 6.9 p.p. among the most seasonal separators, offsetting about one-third of lost earnings.

[22]Seasonal separators also appear to receive more UI benefits than the least seasonal separators during their initial separation, possibly indicating that UI is an implicit part of the labor contract in seasonal sectors (Feldstein, 1975).

In contrast, Appendix Table F.6 shows that changes in means-tested cash-based government transfers are essentially equal to zero for both groups. We do find a modest, marginally significant 0.6 p.p. increase in SNAP benefit receipt among the most seasonal households, relative to a base receipt rate of 10.1 percent for this group. Overall, though, the patterns we observe point to UI providing the main source of social insurance for seasonal separations.

Putting it all together

To tie together the results of the previous analyses, we decompose the response of household income along six mutually exclusive margins: (i) own earnings, (ii) partner earnings, (iii) other co-resident earnings, (iv) household UI income, (v) other household transfer income, and (vi) all other sources. By normalizing each of these components relative to pre-separation earnings, we preserve the additive structure of this decomposition, allowing us to assess how each of these categories contributes to changes in household income.

Table 3 shows the full decomposition of the change in household income between 10 and 13 months after an initial separation for the most seasonal separators. Redisplaying earlier results, we find that own earnings fall among the most seasonal separators and these lost earnings largely pass through to lower household income. Greater transfer income from UI helps increase household income, but this is partly counteracted by lower partner earnings. Other categories have comparatively little influence on the bottom line. In contrast, the least seasonal separators see increases or at most small declines in all components of household income (see Appendix Table F.7).

All told, our preferred pass-through estimate—which incorporates our adjustment for the underreporting of UI receipt—is that household income falls by \$0.81 for each \$1 seasonal loss in earnings. If anything, this estimate is a lower bound. Though we have limited our sample to separations characteristic of seasonal work interruptions, we have presumably misclassified some non-seasonal separators as seasonal. As detailed in Appendix E, our empirical results imply that such misclassification biases our pass-through estimate downward, so that true pass-through is higher.

On net, then, we see little adaptation by seasonal separators to their lost earnings. A possible explanation is that many of the margins along which job separators can usually recoup lost income may be uniquely ill-suited to seasonal separations. During a seasonal dip in the labor market, it will be relatively difficult to find a new job or to have a partner go back to work. At the same time, many seasonal separators will have earned enough to qualify for UI benefits, and so it should not be surprising that UI is the main margin along which they do smooth income.

Differences in adaptation across separators

Seasonal separators may differ in their propensity to recoup lost earnings and in the margins along which they do so. To explore such heterogeneity, and to see whether a particular subgroup of separators is driving our results, we repeat our additive decomposition of household income separately for three different cuts of the sample: splitting (i) by gender, (ii) by labor force status after the initial separation, and (iii) by initial sector (educational services or other). Table 3 shows the decomposition for the most seasonal separators in each subgroup.[23] As before, all changes are normalized by the focal separator's pre-separation earnings, and we focus on estimates adjusted for the underreporting of UI benefits.

Male and female separators exhibit similar declines in both earnings and income. Given previous research documenting the tendency for married women to return to work to smooth shocks to their husbands' income, one might have expected to see an added worker effect for male separators, but we instead find declines in partner earnings for both groups. Male separators receive more UI benefits, but women see relatively more earnings from non-partner co-residents, yielding similar rates of pass-through for men (78%) and women (84%).

We observe larger differences for separations into unemployment versus non-participation. One might have expected separations into unemployment to lead to smaller drops in own

[23] Appendix Table F.7 shows the same set of estimates for the least seasonal separators.

earnings, since many unemployed workers are actively searching for a replacement job, but in fact the decline in earnings is twice as large among the unemployed. However, UI benefits replace a considerably larger share of lost earnings for unemployed workers than for those who initially separated into non-participation, and differences in other income components are roughly offsetting. As a result, separations into unemployment see about one quarter of lost earnings replaced (pass-through of 75%), while separations into non-participation are not insured at all on net (104%).

We next compare separations from educational services versus other sectors. Seasonality in the education sector is distinct in several respects: it reflects policy choices rather than market forces; many school employees are tenured or covered by union contracts; and school staff are often statutorily ineligible for UI benefits. While separators from education indeed show essentially no income recovery via UI benefits, they see an offsetting increase in earnings among non-partner co-residents.[24] On net, pass-through is similar for separations from education (83%) and those from other sectors (80%).

Across all subgroups, lost seasonal earnings pass through to lower household income at high rates. Even after adjusting for underreporting of UI income, each subgroup experiences a loss of household income of at least \$0.75 for every \$1 of lost seasonal earnings. These consistently high pass-through rates indicate that an inability (or unwillingness) to make up for lost earnings is widespread among seasonal workers.

Conclusion

Economic activity fluctuates over the course of the year, causing many workers to experience seasonal work interruptions. The seasonal volatility of employment is echoed in the earnings and incomes of households containing seasonal workers.

[24] Non-partner co-resident earnings increase for seasonal separations from education, but they decrease for seasonal separations from other sectors. A possible explanation for this pattern is that seasonal separations from education tend to occur over the summer, when jobs are relatively plentiful and a separator's teenage or college-age children may be temporarily available for work, whereas other seasonal separations are more concentrated in the winter, when jobs are scarcer and these children are in school.

In this paper, we show how households adapt (and don't adapt) to the income fluctuations created by seasonal work. Seasonal separations are characterized by relatively swift earnings recovery abbreviated by yearly transitions back into joblessness. Seasonal workers seldom find alternative employment in other industries to recoup lost earnings during the off-season, perhaps because many expect to be recalled by their former employers once the off-season ends, or perhaps because jobs are hard to come by during seasonal troughs. We find a modest increase in government transfers in response to a seasonal separation, entirely driven by unemployment insurance, but this is counteracted by contemporaneous declines in partner earnings, a pattern we dub a "subtracted worker effect". On net, households engage in little income-smoothing over the seasonal cycle: each dollar in earnings foregone during the off-season passes through to \$0.81 in reduced household income.

Of course, households may still smooth consumption either by borrowing or by drawing down savings (Morduch, 1995). If households are liquidity constrained or present-biased, however, then seasonal reductions in income may have important implications for consumption volatility (Deaton, 1991; Carroll, 1997; Laibson, 1997). Previous work has shown that the timing of income passes through to consumption at high frequencies, even for sources of income whose arrival dates are entirely predictable, such as paycheck or transfer receipt (Stephens, 2003, 2006; Shapiro, 2005) or the exhaustion of unemployment benefits (Ganong and Noel, 2019; Gerard and Naritomi, 2021). Though our data preclude conclusions about consumption behavior, prior evidence that consumption commonly tracks income suggests that households may find it difficult to maintain constant expenditures through seasonal ups and downs. This possibility is especially relevant because seasonal work is disproportionately common among lower-skilled workers, who are more likely to be liquidity constrained (Jappelli, 1990; Kaplan, Violante, and Weidner, 2014).

The episodic nature of seasonal work may have important ramifications for the design of the social safety net. First, some policies do not readily accommodate workers who deviate from full-year employment. For instance, proposed work requirements for SNAP and

Medicaid would limit eligibility to workers who maintain sufficient employment each month, which could result in seasonal workers losing their eligibility during the off-season (Bauer, Schanzenbach, and Shambaugh, 2018). Second, transfer policies may not be disbursing benefits during the portion of the year when seasonal workers are most in need of assistance. Tax credits like the EITC are typically rebated annually in a single lump-sum payment sometime during the spring months. Changing the timing of these payments to when seasonal workers are typically unemployed could help replace lost income during lean periods and make it easier for families to maintain steady levels of consumption.

References

Barsky, Robert B. and Jeffrey A. Miron. 1989. "The Seasonal Cycle and the Business Cycle." *Journal of Political Economy*, 97(3): 503–534.

Bauer, Lauren, Diane Whitmore Schanzenbach, and Jay Shambaugh. 2018. "Work Requirements and Safety Net Programs." *Economic Analysis*. The Hamilton Project, Brookings Institution, Washington, DC.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2018. "Children, Time Allocation, and Consumption Insurance." *Journal of Political Economy*, 126(S1): S73–S115.

Breiman, Leo. 1997. "Arcing the Edge." Statistics Department, University of California at Berkeley 486.

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics*, 112(1): 1–55.

Davis, Steven J. and Till von Wachter. 2011. "Recessions and the Cost of Job Loss." *Brookings Papers on Economic Activity*, Fall,: 1–72.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica*, 59(5): 1221–1248.

DeBacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos. 2013. "Rising Inequality: Transitory or Persistent? New Evidence from a Panel of US Tax Returns." *Brookings Papers on Economic Activity*, 2013(1): 67–142.

Del Bono, Emilia and Andrea Weber. 2008. "Do Wages Compensate for Anticipated Working Time Restrictions? Evidence from Seasonal Employment in Austria." *Journal of Labor Economics*, 26(1): 181–221.

de Raaf, Shawn, Costa Kapsalis, and Carole Vincent. 2003. "Seasonal Work and Employment Insurance Use." *Perspectives on Labour and Income*, 4(9): 5–11.

Dynarski, Susan and Jonathan Gruber. 1997. "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity*, 1997(1): 229–284.

Faberman, R Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2022. "Job Search Behavior Among the Employed and Non-employed." *Econometrica*, 90(4): 1743–1779.

Farrell, Diana and Fiona Greig. 2015. "Weathering Volatility: Big Data on the Financial Ups and Downs of US Individuals." Washington, DC: JPMorgan Chase Institute.

Feldstein, Martin. 1978. "The Effect of Unemployment Insurance on Temporary Layoff Unemployment." *The American Economic Review*, 68(5): 834–846.

Feldstein, Martin S. 1975. "The Importance of Temporary Layoffs: An Empirical Analysis." *Brookings Papers on Economic Activity*, 1975(3): 725–745.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. 2018. "Integrated Public Use Microdata Series, Current Population Survey: Version 6.0." Minneapolis, MN: IPUMS.

Friedman, Jerome H. 2001. "Greedy Function Approximation: a Gradient Boosting Machine." *Annals of Statistics,*, 1189–1232.

Friedman, Jerome H. 2002. "Stochastic Gradient Boosting." *Computational Statistics & Data Analysis*, 38(4): 367–378.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2001. *The Elements of Statistical Learning.* Vol. First edition, New York, NY: Springer Series in Statistics.

Fromm, Gary. 1979. "Comments on 'An Overview of the Objectives and Framework of Seasonal Adjustment' by Shirley Kallek." In *Seasonal Analysis of Economic Time Series.* , ed. Arnold Zellner, 30–32. Cambridge, MA: National Bureau of Economic Research.

Fujita, Shigeru and Giuseppe Moscarini. 2018. "Recall and Unemployment." *American Economic Review*, 107(12): 3875–3916.

Ganong, Peter and Pascal Noel. 2019. "Consumer Spending During Unemployment: Positive and Normative Implications." *American Economic Review*, 109(7): 2383–2424.

Gerard, François and Joana Naritomi. 2021. "Job Displacement Insurance and (the Lack of) Consumption-Smoothing." *American Economic Review*, 111(3): 899–942.

Geremew, Menelik and François Gourio. 2018. "Seasonal and Business Cycles of US Employment." *Economic Perspectives*, 42(3): 1–28.

Gray, David M. and J. Ted McDonald. 2010. "Seasonal Employment in Canada: Its Decline and Persistence." *Canadian Public Policy*, 36(1): 1–27.

Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review*, 87(1): 192–205.

Hannagan, Anthony and Jonathan Morduch. 2015. "Income Gains and Month-to-Month Income Volatility: Household Evidence from the US Financial Diaries." NYU Wagner Research Paper No. 2659883.

Hyatt, Henry R. 2019. "Co-Working Couples and the Similar Jobs of Dual-Earner Households." *Monthly Labor Review.*

Jacobson, Louis S., Robert J. Lalonde, and Daniel G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review*, 83(4): 685–709.

Jappelli, Tullio. 1990. "Who Is Credit Constrained in the US Economy?" *The Quarterly Journal of Economics*, 105(1): 219–234.

Kalton, Graham, Marianne Winglee, Louis Rizzo, Thomas Jabine, and Daniel Levine. 1998. "SIPP Quality Profile 1998." US Census Bureau SIPP Working Paper 230.

Kaplan, Greg, Giovanni L Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*, 2014(1): 77–153.

Katz, Lawrence F. 1986. "Layoffs, Recalls, and the Duration of Unemployment." NBER Working Paper 1825.

Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury. 2020. "Sources of Displaced Workers' Long-Term Earnings Losses." *American Economic Review*, 110(10): 3231–3266.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 112(2): 443–477.

Lundberg, Shelly. 1985. "The Added Worker Effect." *Journal of Labor Economics*, 3(1, Part 1): 11–37.

Mason, Llew, Jonathan Baxter, Peter L. Bartlett, and Marcus R. Frean. 2000. "Boosting Algorithms As Gradient Descent." 512–518. In *Advances in Neural Information Processing Systems.*

Meyer, Bruce D., Wallace K.C. Mok, and James X. Sullivan. 2015. "The Under-Reporting of Transfers in Household Surveys: Its Nature and Consequences." Working paper (previously issued as NBER Working Paper 15181).

Morduch, Jonathan. 1995. "Income Smoothing and Consumption Smoothing." *Journal of Economic Perspectives*, 9(3): 103–114.

Moretti, Enrico. 2000. "Do Wages Compensate for Risk of Unemployment? Parametric and Semiparametric Evidence from Seasonal Jobs." *Journal of Risk and Uncertainty*, 20(1): 45–66.

Morris, Pamela, Heather Hill, Lisa A. Gennetian, Chris Rodrigues, and Sharon Wolf. 2015. *Income Volatility in US Households with Children: Another Growing Disparity Between the Rich and the Poor?* Madison, WI:University of Wisconsin-Madison, Institute for Research on Poverty.

Nekoei, Arash and Andrea Weber. 2015. "Recall Expectations and Duration Dependence." *American Economic Review*, 105(5): 142–146.

Nekoei, Arash and Andrea Weber. 2020. "Seven Facts About Temporary Layoffs." CEPR Discussion Paper 14845.

Ngai, L. Rachel and Silvana Tenreyro. 2014. "Hot and Cold Seasons in the Housing Market." *American Economic Review*, 104(12): 3991–4026.

Olivei, Giovanni and Silvana Tenreyro. 2007. "The Timing of Monetary Policy Shocks." *American Economic Review*, 97(3): 636–663.

Price, Brendan M and Melanie Wasserman. 2024. "The Summer Drop in Female Employment." *Review of Economics and Statistics*, 1–46.

Shapiro, Jesse M. 2005. "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle." *Journal of Public Economics*, 89(2–3): 303–325.

Sharpe, Andrew and Jeremy Smith. 2005. "Labour Market Seasonality in Canada: Trends and Policy Institutions." Centre for the Study of Living Standards. CSLS Research Report Number 2005-01.

Stephens, Melvin Jr. 2002. "Worker Displacement and the Added Worker Effect." *Journal of Labor Economics*, 20(3): 504–537.

Stephens, Melvin Jr. 2003. "'3rd of tha Month': Do Social Security Recipients Smooth Consumption Between Checks?" *American Economic Review*, 93(1): 406–422.

Stephens, Melvin Jr. 2006. "Paycheque Receipt and the Timing of Consumption." *Economic Journal*, 116(513): 680–701.

Stinson, Martha. 2003. "Technical Description of SIPP Job Identification Number Editing in the 1990–1993 SIPP Panels." Unpublished manuscript, US Census Bureau.

Warner, Elizabeth J. and Robert B. Barsky. 1995. "The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays." *The Quarterly Journal of Economics*, 110(2): 321–352.

Figure 1: Non–seasonally adjusted unemployment among prime-age US workers

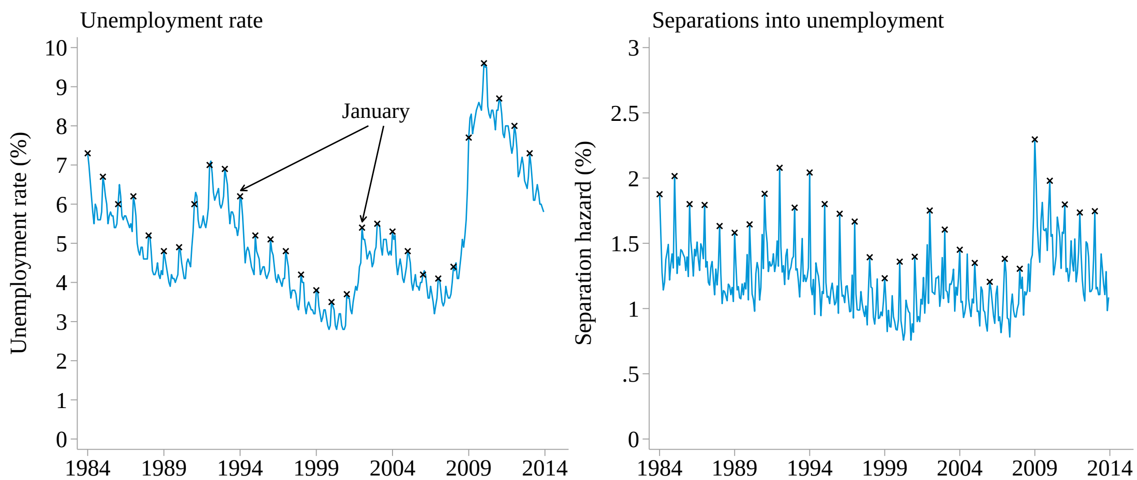

Notes: The left panel plots the raw monthly unemployment rate among US workers ages 25–54, as reported by the Bureau of Labor Statistics (series LNU04000060). The right panel plots, using the Current Population Survey, the monthly hazard rate of separating from a job into unemployment (i.e., the probability of being employed during the reference week in month $t-1$ and unemployed during the reference week in month t). We impute the hazard rate for eight observations (none in January) in which fewer than 90% of workers can be matched to the prior month.

Figure 2: Excess recurrence of job separations at annual frequencies

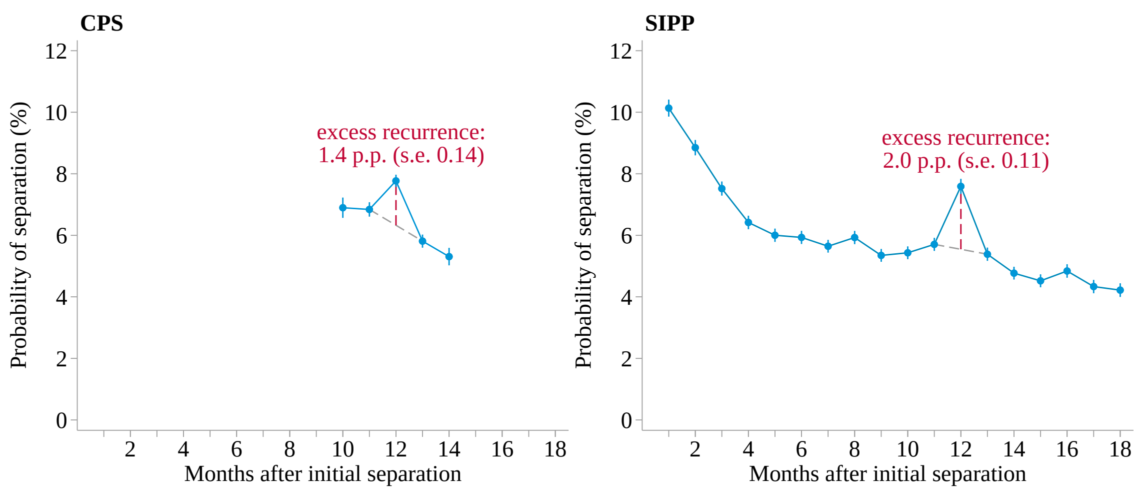

Notes: Probabilities $\hat{\rho}_\tau$ of experiencing a recurrent separation τ months after an initial separation from employment into non-employment, estimated among prime-age CPS and SIPP separators. Estimates are obtained via regression as described in the text. Spikes show 95% confidence intervals, clustered by individual.

Figure 3: Share of construction separators classified as "most seasonal"

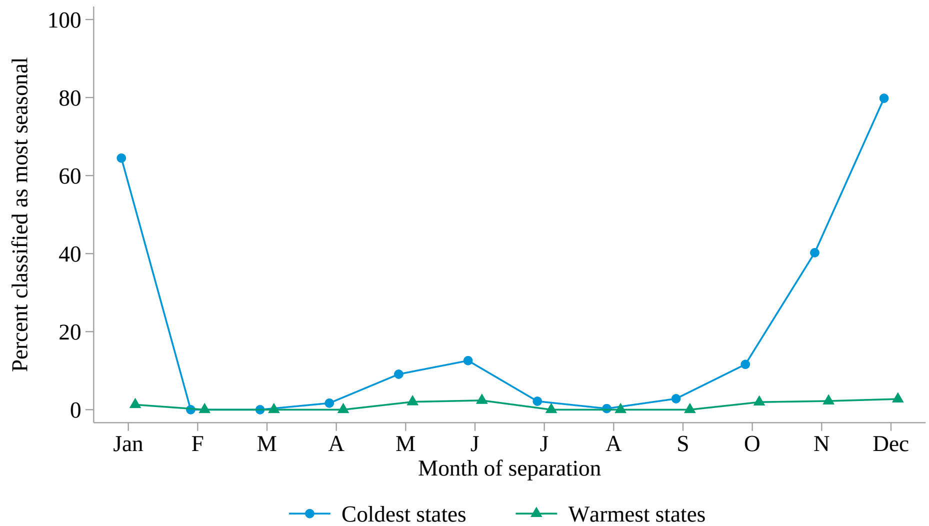

Notes: Share of SIPP separations from construction-sector jobs that fall within the top decile of predicted excess recurrence, according to our predictive algorithm. The "coldest" [respectively, warmest] states are those that fall within the lowest [highest] tercile of average January maximum temperature over 1984–2013.

Figure 4: Earnings and household income of most/least-seasonal separators

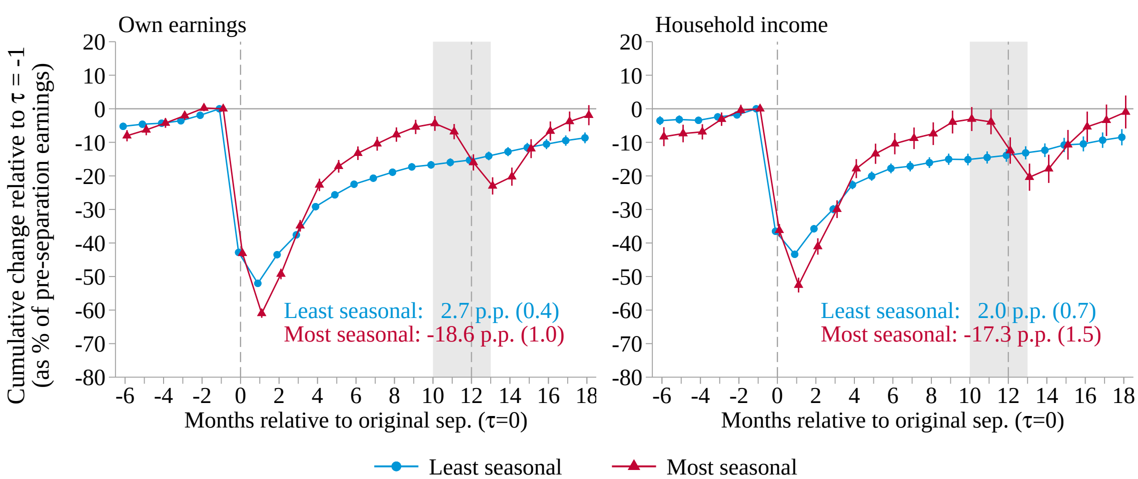

Notes: Left panel plots the evolution of personal earnings for prime-age SIPP respondents who experience an initial separation and are in either the top decile ("most seasonal") or bottom half ("least seasonal") of predicted excess recurrence. Right panel plots household income for this same sample. Both variables are expressed as a percentage of the individual's earnings one month prior to the initial separation. Spikes show 95% confidence intervals, clustered by individual.

Figure 5: Changes in partner earnings, decomposed by industry

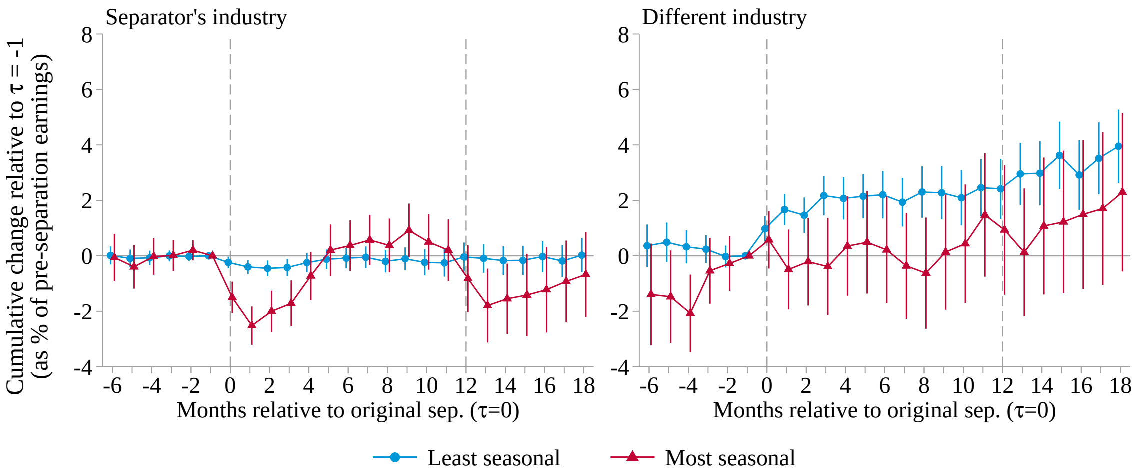

Notes: Changes in partner's earnings by industry for SIPP separators in the top decile ("most seasonal") and bottom half ("least seasonal") of the predicted excess recurrence distribution. Partner's earnings are divided into earnings from jobs in the same industry as the separator's job one month before the *initial* separation, and earnings from jobs in all other industries. All variables are expressed as percentages of the separator's earnings one month before the initial separation. Spikes show 95% confidence intervals, clustered by individual.

Figure 6: Changes in receipt of unemployment benefits, adjusted for underreporting

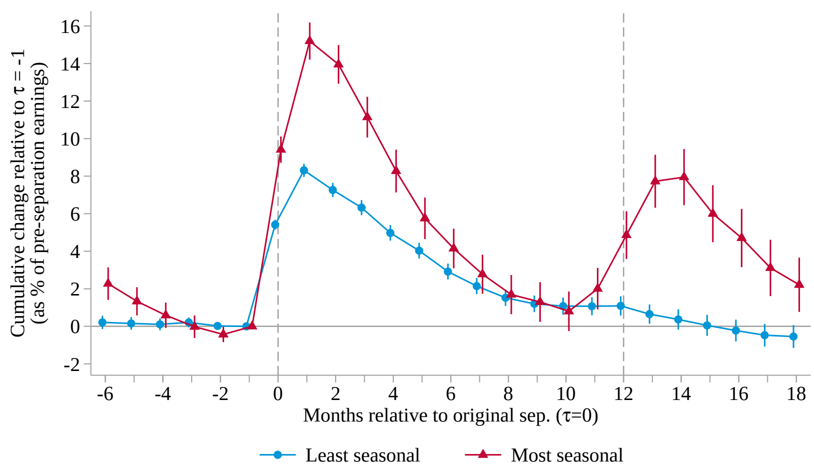

Notes: Changes in household UI income for SIPP separators in the top decile ("most seasonal") and bottom half ("least seasonal") of the predicted excess recurrence distribution. We scale up UI income at the individual level using the year-by-year estimates reported in Meyer, Mok, and Sullivan (2015) to adjust for underreporting of UI transfers in survey data relative to administrative totals. UI income is expressed as a percentage of the separator's earnings one month preceding the initial separation. Spikes show 95% confidence intervals, clustered by individual.

Table 1: Summary statistics: SIPP separators least/most likely to separate again

	Least seasonal (N = 61,274)		Most seasonal (N = 12,678)	
Demographics				
Female	48.1	(50.0)	47.8	(50.0)
Age	37.0	(8.3)	37.6	(8.4)
Non-white	30.5	(46.0)	26.8	(44.3)
Non-college	52.0	(50.0)	51.6	(50.0)
Household structure				
Spouse or partner in household	60.2	(48.9)	65.4	(47.6)
Spouse/partner employed (if present)	77.2	(42.0)	78.4	(41.1)
Spouse/partner in same industry (if employed)	18.2	(38.6)	20.8	(40.6)
≥1 children in household	55.5	(49.7)	57.9	(49.4)
Pre-separation industries (select examples)				
Agriculture, fishing, & forestry	0.6	(7.5)	12.0	(32.5)
Construction	12.2	(32.8)	14.7	(35.4)
Educational services	4.9	(21.7)	27.6	(44.7)
Healthcare	11.3	(31.7)	0.3	(5.5)
Pre-separation monthly earnings (2017$)				
Personal earnings	2,678.4	(2862.3)	2,532.6	(2192.4)
Household income	5,709.4	(4960.9)	5,838.0	(4535.9)
Separated into …				
Separated into unemployment	58.6	(49.3)	64.3	(47.9)
Separated into non-participation	41.4	(49.3)	35.7	(47.9)

Notes: All columns restrict to job separators ages 25–54 observed during 1984–2013. "Least seasonal": restrict to separators in the bottom half of predicted excess recurrence. "Most seasonal": restrict to separators in the top decile of predicted excess recurrence. Separators' industries, earnings, and income are measured one month prior to separation. Other than age, earnings, and income, all statistics are expressed as percentages (s.d. in parentheses). Decile cutoffs are calculated using sampling weights, so sample sizes differ between deciles.

Table 2: Changes in separator earnings by exposure to seasonal work interruptions

All SIPP panels	Most seasonal	Least seasonal	**1990–1993 panels only**	Most seasonal	Least seasonal
Total earnings	-18.62	2.67	Total earnings	-24.92	3.00
	(1.04)	(0.42)		(2.09)	(0.76)
Original industry	-19.80	-0.54	*Original employer*	-22.69	-1.96
	(0.90)	(0.31)		(1.68)	(0.50)
			Other employers in original industry	-2.82	1.23
				(1.12)	(0.45)
Other industries	1.18	3.21	*Other industries*	0.58	3.72
	(0.61)	(0.33)		(1.03)	(0.58)

Notes: Changes in components of separators' earnings from 10 to 13 months after an initial separation for SIPP separators in the top decile ("most seasonal") and bottom half ("least seasonal") of the predicted excess recurrence distribution. The right panel uses the 1990–1993 SIPP sample with corrected job identifiers (Stinson, 2003). All variables are expressed as percentages of the individual's earnings one month prior to the initial separation. Standard errors are shown in parentheses.

Table 3: Decomposing changes in household income among probable seasonal separators

	Overall	Sex of focal separator: Male	Female	Type of separation: Unemp.	Non-part.	Pre-separation sector: Education	Other
A. Unadjusted estimates							
Own earnings	-18.62	-17.81	-19.42	-22.56	-11.22	-22.91	-16.96
	(1.04)	(1.52)	(1.42)	(1.29)	(1.71)	(2.02)	(1.21)
Partner earnings	-2.60	-2.07	-3.15	-1.94	-3.86	-1.13	-3.18
	(0.97)	(0.75)	(1.82)	(0.93)	(2.14)	(2.37)	(1.00)
Other co-resident earnings	-1.13	-3.43	1.23	-1.95	0.42	4.00	-3.04
	(0.81)	(1.10)	(1.18)	(0.94)	(1.48)	(1.55)	(0.94)
Household UI benefits	4.65	6.34	2.91	5.97	2.17	0.19	6.32
	(0.40)	(0.60)	(0.54)	(0.52)	(0.59)	(0.40)	(0.53)
Other transfer receipts	0.19	0.19	0.20	0.13	0.31	0.10	0.23
	(0.13)	(0.20)	(0.17)	(0.18)	(0.20)	(0.11)	(0.18)
All other sources	0.16	-0.05	0.39	0.61	-0.68	0.54	0.03
	(0.50)	(0.65)	(0.76)	(0.56)	(0.98)	(1.17)	(0.53)
Household income	-17.34	-16.83	-17.83	-19.75	-12.87	-19.22	-16.61
	(1.55)	(1.88)	(2.48)	(1.76)	(2.97)	(3.45)	(1.70)
Pass-through rate	93.1	94.5	91.8	87.5	114.7	83.9	97.9
	(6.8)	(8.0)	(11.0)	(6.1)	(23.1)	(12.5)	(8.3)
B. Adj. for UI underreporting							
Household UI benefits	6.93	9.29	4.49	8.80	3.40	0.34	9.39
	(0.60)	(0.88)	(0.80)	(0.76)	(0.88)	(0.57)	(0.78)
Household income	-15.06	-13.87	-16.26	-16.92	-11.64	-19.06	-13.54
	(1.55)	(1.87)	(2.48)	(1.75)	(2.97)	(3.46)	(1.69)
Pass-through rate	80.9	77.9	83.7	75.0	103.7	83.2	79.8
	(7.0)	(8.4)	(11.1)	(6.4)	(23.1)	(12.6)	(8.5)

Notes: Changes in components of household income between 10 and 13 months after an initial separation for SIPP respondents in the top ("most seasonal") decile of the predicted excess recurrence distribution. All variables are expressed as percentages of the individual's earnings one month prior to the initial separation. Standard errors shown in parentheses.

FOR ONLINE PUBLICATION

Appendix A Details on Sample Construction

We use data from three sources: the Current Population Survey (CPS), the Survey of Income and Program Participation (SIPP), and the National Oceanographic and Atmospheric Administration (NOAA).

CPS data

We use monthly core CPS files provided by the Integrated Public Use Microdata Series, or IPUMS (Flood et al., 2018), for the years 1984–2013. Each CPS household is interviewed for four consecutive months, out of rotation for the next eight months, and then interviewed again for a final four months. We link observations both month-to-month and year-over-year using household- and person-level identifiers provided by IPUMS.

We restrict our sample to the civilian population ages 25–54. We weight respondents by the final person-level weights provided by IPUMS. Except where indicated, we employ an unbalanced panel throughout our analysis, with individuals retained in-sample as long as possible to maximize statistical power.

Using the IPUMS variables EDUC, IND1990, and OCC1990 (which have been harmonized over time), we recode education, industry, and occupation to categories compatible with the SIPP. We recode education to the four categories “less than high school”, “high school graduate”, “some college”, and “college graduate” (including anyone with 4+ years of college). We aggregate detailed industry codes into 15 one-digit sectors, augmented with a missing category. We likewise aggregate detailed occupational codes into 14 one-digit occupations plus a missing category.

Separations in the CPS reflect changes in employment status between consecutive months’ reference weeks. Importantly, the length of time between CPS reference weeks varies both month-to-month and year-to-year. In most months, the reference week is the 7-day calendar week (Sunday–Saturday) that contains the 12th day of the month, but it is sometimes shifted to avoid contacting households during holiday periods. According to the Bureau of Labor Statistics, the November reference week is moved one week earlier if Thanksgiving falls during the week that contains the 19th day of the month (and sometimes in other years, at the Census Bureau’s discretion). The December reference week is moved one week earlier if the week that straddles the 5th day of the month is contained entirely within December. We apply these rules to calculate the number of weeks elapsed between the reference weeks in each pair of successive months. Since longer gaps between reference weeks tend to increase the number of observed separations (as workers have more time to separate), our CPS regressions control for our measure of weeks elapsed.

SIPP data

The SIPP is structured as a series of panels, each of which is named after the year in which respondents were initially interviewed. New panels were introduced annually from 1984 to 1988 and from 1990 to 1993, then (after a major redesign) in the years 1996, 2001, 2004, 2008, and 2014. Whereas the 1984–2008 panels interviewed respondents three times per year, the 2014 panel interviews each respondent only once per year. Because these lengthy gaps between interviews limit the usefulness of the 2014 panel for measuring month-to-month employment flows, we confine our analysis to the 1984–2008 panels. Altogether, our data straddle the period 1984–2013.

We begin by assembling raw SIPP extracts published by the National Bureau of Economic Research (NBER).[25] For the pre-1996 panels, the NBER provides separate "wave" files covering each four-month interview window as well as "full-panel" files that stack observations across all waves within a given panel. For these panels, we rely principally on the full-panel files but merge in weekly detail on employment status that is only reported in the wave files. From the 1996 panel onwards, the Census Bureau discontinued the full-panel files, so we rely entirely on the wave files.

From the full universe of SIPP observations, we impose the following sample restrictions. First, we delete observations with missing household identifiers or invalid interview codes, such as cases of "Type Z" person non-response that the Census Bureau populates through imputation. Second, we restrict to ages 25–54. Third, we drop person-months in which an individual belongs to the Armed Forces. Lastly, we discard observations for which the Census Bureau's cross-sectional sampling weights are zero or missing. To adjust for differences in sampling probability and attrition rates, we compute all statistics using these weights, so that our estimates are representative of the US population at each point in time. As with the CPS, we employ an unbalanced SIPP panel throughout our analysis, with individuals retained in-sample for as many observations as possible.

We recode education, industry, and occupation to the same categories described above for our CPS sample. We deflate all monetary amounts to January 2017 dollars using the Federal Reserve Board's chain-type personal consumption expenditures price index.

Although the SIPP collects weekly and monthly information from respondents, the underlying interviews are conducted only once every four months. In each interview session, respondents answer questions about each of the past four months, requiring them to recall information from these prior months. Given the difficulty of remembering the exact timing of employment changes further in the past, the SIPP interview structure leads individuals to underreport changes in their employment status during each four-month interview wave, which results in some transitions that actually occurred *within* an interview period being recorded as happening *between* interview periods (Kalton et al., 1998). This phenomenon—known as "seam bias", in reference to the seams between recall periods—leads to an artificial excess of employment and unemployment spells with durations that are exact multiples of four months. Appendix C reports estimates of annually recurrent work interruptions that correct for the presence of seam bias.

NOAA data

To measure each state's usual January temperature, we compute the average maximum daily temperature for the Januaries of 1984–2013 using state-level data from NOAA's nCLIMDIV

[25]The NBER converts the underlying data from the Census Bureau into Stata and other formats to facilitate ease of access. See `http://data.nber.org/data/survey-of-income-and-program-participation-sipp-data.html`.

database. The nCLIMDIV data do not report temperatures for Hawaii or for the District of Columbia (DC). We set Hawaii's average January temperature using that for Florida (the warmest state with non-missing data), and we set DC's temperature equal to that for Maryland.

To protect respondent confidentiality, the Census Bureau sometimes combines certain less-populous states when reporting state of residence in the SIPP (e.g., binning Maine and Vermont together). In such cases, we assign January temperature using the simple unweighted average of January temperatures in the binned states.

Appendix B Conceptual Framework

To build intuition for our empirical analysis, we sketch a simple model in which a subset of workers are employed in seasonal jobs that end at particular times in the calendar year. In this stylized environment, our approach of measuring the excess mass of annually recurrent job separations allows us to identify, in a probabilistic sense, which work interruptions are seasonal in nature.

Consider an economy with a fraction α of seasonal workers and a fraction $1 - \alpha$ of non-seasonal workers. For expositional simplicity, we adopt the polar assumption that each seasonal worker i has deterministic employment, always working from month m_i to m_i' each year. The timing of the starting and ending months may vary across individuals: for example, construction workers may be employed from April through December, ski instructors from November through May, and school bus drivers from September through June. Each seasonal worker is assumed to remain jobless during the "off-season" and thus experiences exactly one hire and one separation each year. By contrast, each non-seasonal worker has stochastic employment governed by constant one-month job-finding and separation hazard rates f_i and s_i, so that the steady-state employment probability equals $\frac{f_i}{s_i+f_i}$ at all points in the year.

How might we determine the prevalence of seasonal workers in such an economy? One approach would be to tally up the absolute value of the month-to-month changes in *aggregate* employment throughout the year, dividing by two to avoid double-counting (as each seasonal increase must be accompanied by a seasonal decrease). While easily implemented using aggregate data, this approach would yield only a lower bound on α because any cross-worker differences in the timing of the seasonal cycle—as between construction workers and bus drivers—will cancel out in the aggregate. In the extreme, an economy consisting entirely of seasonal workers ($\alpha = 1$) could exhibit zero seasonality in aggregate employment, for example if half of all workers are employed in the first half of the year while the remainder are employed in the latter half of the year.

To accommodate arbitrary heterogeneity in the timing of each worker's seasonal cycle, we propose a method based on the periodicity of individual employment flows. Let $\mathrm{Pr}_i(\mathrm{Sep}_t)$ be the probability that worker i separates from a job in month t, i.e., was employed in $t-1$ but not in t. Among newly separated workers, the probability of experiencing a *recurrent* separation exactly τ months later is given by $\mathrm{Pr}_i(\mathrm{Sep}_{t+\tau} \mid \mathrm{Sep}_t)$. Non-seasonal workers are approximately as likely to separate from a new job either 11, 12, or 13 months after a previous separation.[26] Seasonal workers, on the other hand, will only separate from a job exactly 12 months after a previous separation, so that $\Pr_i(\text{Sep}_{t+\tau} \mid \text{Sep}_t) = \mathbb{1}\{\tau = 12\}$. This reasoning suggests that the share of seasonal workers within a group of new separators can be gleaned from the *excess recurrence* of separations at 12 months, i.e., the recurrence rate for $\tau = 12$ in excess of the average for $\tau = 11$ and $\tau = 13$. Letting $\rho_\tau \equiv \Pr(\text{Sep}_{t+\tau} \mid \text{Sep}_t)$ denote the average recurrence rate in the economy, we define

$$\text{excess recurrence} \equiv \Pr(\text{Sep}_{t+12} \mid \text{Sep}_t) - \frac{1}{2}\Big(\Pr(\text{Sep}_{t+11} \mid \text{Sep}_t) + \Pr(\text{Sep}_{t+13} \mid \text{Sep}_t)\Big) \tag{6}$$

which can be written compactly as: $\rho_{12} - \frac{1}{2}(\rho_{11} + \rho_{13})$.

Although it is not a direct estimate of the fraction of seasonal workers α, excess recurrence can be expressed in terms of α:

$$\text{excess recurrence} \approx \frac{\frac{\alpha}{12}}{(1-\alpha)\int \frac{s_i f_i}{s_i + f_i} di + \frac{\alpha}{12}} \tag{7}$$

where the denominator is equal to the average monthly separation rate in the economy, Pr(Sep).[27]

Excess recurrence tells us the share of *separations* that repeat at annual intervals, rather than the share of *workers* who experience such repeated separations. Multiplying this quantity by the average separation rate and then multiplying by 12 gives the total proportion of seasonal workers in this economy:

$$\text{excess recurrence} \cdot \Pr(\text{Sep}) \cdot 12 \approx \alpha \tag{8}$$

In this way, measuring excess recurrence allows us to pin down the prevalence of seasonal work interruptions in a given population, a statistic that cannot be determined from aggregate data alone. These rescaled estimates are especially helpful when we wish to compare different segments of the labor market with differing degrees of churn (as measured by separations per worker), as we do in Appendix D.

This stylized model makes a number of simplifying assumptions, but excess recurrence is still likely to be informative when these assumptions are relaxed. If some individuals are stochastically seasonal, perhaps likely but not certain to be laid off at the same time every year, then excess recurrence will yield a lower bound on the share of workers prone to seasonal work interruptions. Additionally, if some workers change from being seasonal to non-seasonal throughout their careers and vice versa, then excess recurrence will estimate the share of workers who remain seasonal in consecutive years. This quantity may be of direct interest, but it too would form a lower bound on the share of workers who are seasonal at any point in their life.

[26]For a non-seasonal worker, $\Pr_i(\text{Sep}_{t+\tau} \mid \text{Sep}_t) = \frac{s_i f_i}{s_i+f_i}(1-(1-s_i-f_i)^{\tau-1})$, which converges to $\frac{s_i f_i}{s_i+f_i}$ as τ becomes large. For empirically plausible job-finding and separation rates, this probability is close to constant for $\tau \in \{11, 12, 13\}$.

[27]The approximation error in this expression arises because, for generic values of s_i and f_i, the separation rate ρ_τ among non-seasonal separators will not yet have converged to its steady-state value by $\tau = 11$ or $\tau = 13$. For empirically relevant hazard rates, however, the approximation error is negligible.

Beyond revealing the prevalence of seasonal work in the labor market as a whole, excess recurrence can also tell us about the kinds of workers employed in seasonal jobs. To illustrate this point, suppose that the labor force can be partitioned into G groups indexed by g (e.g., education groups). By computing excess recurrence separately for each group and scaling it by the group's monthly separation rate, we can estimate the share of seasonal workers α_g in each group. By comparing α_g across groups, we can gauge how seasonal work is distributed throughout the economy, affording us a window into seasonal cycles unobscured by any offsetting seasonal cycles that exist in aggregated data. Moreover, to the extent that annually recurrent separations can be predicted on the basis of observable individual characteristics, our method points the way towards identifying individual workers who are likely to be seasonally employed.

Appendix C Robustness of Excess Recurrence

Our CPS and SIPP estimates of excess recurrence survive a suite of sensitivity checks, which confirm that our measure of seasonal work interruptions rests on a robust empirical fact. We present these robustness checks in Appendix Table F.3.

Month duration and seam bias In Figure 2, some of the excess separations measured at $\tau = 12$ may be unrelated to seasonality, but rather reflect mechanical factors that induce spurious autocorrelation at annual frequencies.

First, in both the CPS and the SIPP, some monthly observations encompass a larger number of weeks than others, owing to the particular timing of days and holidays within the month. Because months with a larger number of days tend to be coded as spanning a larger number of weeks, because month duration is (almost) constant across years, and because a larger number of weeks affords more opportunities for a work interruption, the raw recurrence probabilities at $\tau = 12$ may be biased upward.

Second, as explained in Appendix A, the SIPP's interview structure gives rise to a pattern of "seam bias" whereby a disproportionate number of employment transitions are observed at the seams between four-month interview periods. Since a disproportionate share of "initial" job separations coincide with interview seams, subsequent observations at multiples of four months also tend to align with interview seams. Seam bias is evident in the elevated separation probabilities at months 8 and 16 compared to nearby months, and thus it likely explains some of the elevation at month 12 as well.[28]

Appendix Figure F.1 reports regression estimates that control for these potential sources of bias. First, for both the CPS and the SIPP we control for the number of weeks represented by each monthly observation. Second, to account for the possibility of seam bias, the SIPP specification includes indicator variables for person-months occurring immediately after an interview seam. Since the magnitude of seam bias appears to vary between an individual's *first* and *subsequent* post-baseline interview seams, we include one indicator variable

[28] While Figure 2 reveals no similarly elevated separation probability 4 months after baseline, the underlying week-level data shows a less pronounced, but qualitatively similar spike occurring at the 4 month seam as well. The muted spike at this horizon may reflect the mechanical linkage between separations at $\tau = 0$ and separations at $\tau = 4$ (which for many respondents draw on information from the same interview wave).

for each respondent's first monthly observation that coincides with a seam, and another indicator variable for all subsequent observations that coincide with seams.

In the CPS, our estimate of excess recurrence is little changed at 1.5 p.p. (s.e. = 0.14). In the SIPP, incorporating these controls eliminates the excess separations observed at $\tau = 8$ and 16, and it yields a slightly attenuated estimate of excess recurrence at 1.6 p.p. (s.e. = 0.11). We include these control variables in our estimates of excess recurrence in Appendix D.

Sample attrition To maximize statistical power, our SIPP specifications retain all available observations in the 18 months following each initial separation, even if an individual exits from the sample before these 18 months are up. As such, the group of separators used in estimation varies with the horizon τ. Respondents can disappear from the SIPP either because their survey panel officially comes to an end—which is unlikely to bias our estimates—or because of survey attrition, which is potentially non-random.

To assess whether dynamic changes in sample composition are biasing our estimates, we re-estimate our preferred specification using a balanced panel of SIPP respondents who are continuously present from month $t_0 - 1$ through month $t_0 + 18$. We obtain similar point estimates in this balanced sample, with excess recurrence estimated to equal 1.5 p.p. (s.e. = 0.13), indicating that attrition does not meaningfully bias our estimates of ρ_τ.

Changes over time Next, we verify that our estimates are not driven by changes in the SIPP survey over time. The design of the SIPP was changed substantially for the 1996 panel, including increasing the overall size and duration of panels, eliminating overlapping panels, and using computer-assisted interviewing. As a result, the post-1996 panels may have different patterns of seam bias than preceding panels. We estimate the excess recurrence rate separately for the pre-1996 and post-1996 panels, obtaining estimates of 1.7 p.p. (s.e. = 0.16) and 1.5 p.p. (s.e. = 0.13), respectively.

Excess recurrence in job-finding As a final check, we note that seasonality should manifest in annually recurrent *job-finding* as well as annually recurrent *separations*. To check this intuition, we compute the probability of a recurrent transition from non-employment into employment among SIPP respondents who make such a transition in a baseline period. The resulting profile of recurrent transitions into work is very similar to that for recurrent separations, lending further credence to the association between employment seasonality and the annual periodicity in labor market flows. The excess recurrence of job-finding at a 12-month frequency equals 1.5 p.p. (s.e. = 0.10), almost identical to the annual excess recurrence of job separations.

Appendix D Connecting Recurrent Separations to Labor Market Seasonality

Our proxy for seasonal work interruptions—the excess recurrence of separations at annual frequencies—aligns with three characteristic features of labor market seasonality. Excess recurrence is especially pronounced (i) at times of year when aggregate employment is declining, (ii) in sectors subject to large seasonal swings in employment, and (iii) among job losers on temporary layoff. Though nothing in our method creates a mechanical linkage between periodicity in individual separations and seasonal patterns in the labor market, we find close connections between these two phenomena.

Timing within the year Net of secular trends, the employment-to-population rate among prime-age workers typically declines both at the start of winter (likely reflecting adverse weather) and at the start of summer (likely reflecting schools' summer breaks).[29] If annually recurrent separations stem from the same seasonal forces that operate in the aggregate, then one might expect such separations to coincide with the aggregate cycle.

Appendix Figure F.2 uses our CPS sample to estimate excess recurrence separately by the month in which the original separation occurred. Consistent with the seasonal timing of declines in aggregate employment, we find that individuals who experience separations in January and June are the most (excessively) likely to separate again 12 months later.[30]

Heterogeneity across industries The left panel of Appendix Figure F.3 estimates excess recurrence in the CPS separately by pre-separation industry, using 15 one-digit sectors. The three leading sectors—(i) agriculture, fishing, and forestry, (ii) entertainment and recreation, and (iii) educational services—line up neatly with the seasonal rhythms in farming, leisure, and schooling.[31] By contrast, the health care sector—in which employment fluctuates little throughout the year—shows no excess recurrence of individual separations. Most sectors exhibit some degree of excess recurrence, however, indicating that seasonality is somewhat widespread and not exclusively confined to a few small sectors.

Excess recurrence tells us the share of *separations* that repeat at annual intervals, rather than the share of *workers* who experience such repeated separations. As detailed in Appendix B, multiplying excess recurrence by the monthly separation rate in a given industry captures the probability that a worker in that industry will experience an annually

[29] In a regression that uses a flexible time trend to control for secular and business-cycle fluctuations, the employment-to-population rate among prime-age CPS respondents declines in November (0.1 p.p.), December (0.3 p.p.), January (0.8 p.p.), June (0.3 p.p.), and July (0.2 p.p.).

[30] Although the CPS reference week normally encompasses the 12th day of the month, the December reference week is usually advanced by one week so that interviews conclude before Christmas. As a result, a disproportionate share of December separations first appear in the CPS in January.

[31] Interestingly, excess recurrence is relatively low in the retail sector, which appears quite seasonal by other metrics. This may reflect the brevity of the holiday shopping season: since retail's "off-season" lasts for most of the year, workers laid off after Christmas may seek work in other sectors rather than waiting around for the next holiday boom.

recurrent separation. The right panel of Appendix Figure F.3 presents estimates of

$$\hat{\alpha}_i \equiv \text{excess recurrence in industry } i \cdot \widehat{\Pr}(\text{Sep} \mid \text{industry} = i) \cdot 12 \tag{9}$$

Scaling excess recurrence by the monthly separation rate widens the gap between agriculture and the other industries: seasonal work interruptions are most common in agriculture both because many workers separate from agriculture in the average month and because seasonal workers represent a large share of separators in agriculture. Scaling by the separation rate also makes the construction industry stand out from industries with similar rates of excess recurrence, indicating that construction exhibits a large number of both seasonal and non-seasonal work interruptions.

Connections to temporary layoffs Many job losses turn out to be temporary layoffs from which workers are eventually recalled to their previous employers. While not all seasonal separations result from job loss, layoffs induced by seasonal declines in labor demand are likely to be temporary because a seasonal employer's "off-season" is generally both short-lived and predictable in duration. Intuitively, workers who are laid off at the end of the harvest cycle, construction schedule, tourist season, or school year seem likely to be rehired when demand picks back up.

To explore this idea, we leverage the fact that the CPS asks unemployed workers why they are unemployed. Appendix Figure F.4 reports estimates of excess recurrence for workers whose initial separation was into unemployment, first for the full subsample and then stratifying by reason for unemployment. In each case, we compute the excess probability that a worker experiences a second separation *of any kind*, including separations of types other than the first separation. For unemployed workers as a whole, we estimate excess recurrence to be 1.6 p.p., similar to our baseline estimates. Within this group, annually recurrent separations are driven primarily by workers reporting that they are on temporary layoff (estimate = 3.9 p.p., s.e. = 0.44), meaning that they expect to return to their previous employer within 6 months or have been given a specific recall date. We find small and statistically insignificant estimates for the remaining categories.

To the extent that employers in seasonal industries engage in temporary layoffs, we might expect to observe repeat separations from the same employer staggered at 12-month intervals. The 1990–1993 SIPP panels include reliable employer and industry identifiers (Stinson, 2003), which allow us to partition recurrent separations into departures from the *same employer* from which the worker separated the first time, separations from other employers in the *same industry*, and separations from *different industries*.[32] As shown in Appendix Figure F.5, the spike in job-finding at annual frequencies is primarily accounted for by repeat separations from the same employer. This result, which confirms a similar finding for Austria in Del Bono and Weber (2008), draws a close empirical connection between the

[32]To accommodate dual job-holding, the SIPP records up to two employer IDs in each month. We code the second separation as pertaining to the same employer if the recurrent separation aligns with either of the two employer IDs associated with the first separation. In cases where employer IDs are missing, we presume that workers have separated from distinct employers. In cases where industry identifiers are missing, we likewise presume that workers have separated from distinct industries. These conservative coding choices will, if anything, lead us to understate the same-employer component of excess recurrence.

periodic separations identified by our method and the layoff-recall phenomenon studied by prior literature.[33]

Appendix E Bounding the Pass-through Rate

Formally, let $\delta_{i,t}$ be an indicator for an individual i experiencing a seasonal work interruption in time period t. We are interested in the quantities

$$\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \delta_{i,t}=1\right], \quad \mathbb{E}\left[\Delta\text{Income}_{i,t} \mid \delta_{i,t}=1\right], \quad \frac{\mathbb{E}\left[\Delta\text{Income}_{i,t} \mid \delta_{i,t}=1\right]}{\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \delta_{i,t}=1\right]}, \tag{10}$$

with the differences computed in some window around period t (which may vary across individuals). If $\delta_{i,t}$ were directly observable for each individual in our dataset, estimating these quantities would be as simple as plugging in the sample analogues for each expectation. As we noted earlier, however, household surveys do not typically ask about seasonal work status; moreover, it is unclear how self-reported seasonal status would relate to seasonality as understood by economists.

Instead, we use our measure of excess recurrence to construct proxies for seasonal work interruptions. With a proxy indicator $\hat{\delta}_{i,t}$ for seasonal work interruption at date t, our key assumption is that changes in earnings around date t among individuals tagged by this indicator approximate our first object of interest:

$$\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \hat{\delta}_{i,t}=1\right] \approx \mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \delta_{i,t}=1\right] \tag{11}$$

We can compute approximations for the other objects of interest in the same fashion, by conditioning the expectations on the proxy indicator $\hat{\delta}_{i,t}$ in place of the unobserved indicator of seasonal work interruptions $\delta_{i,t}$.

To differing degrees, both of our proxy indicators may include some non-seasonal workers by chance, which can attenuate our estimates of the changes in earnings and household income. Suppose that among observations for which our proxy indicates a seasonal work interruption (i.e., $\hat{\delta}_{i,t}=1$), a fraction p of these observations represent actual seasonal work interruptions ($\delta_{i,t}=1$), while a fraction $(1-p)$ represent non-seasonal workers ($\delta_{i,t}=0$). Our estimate of the change in earnings will be a mixture of the changes for the two groups,

$$\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \hat{\delta}_{i,t}=1\right] = p\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \delta_{i,t}=1\right] + (1-p)\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \delta_{i,t}=0\right] \tag{12}$$

with the same identity also holding with $\text{Income}_{i,t}$ in place of $\text{Earnings}_{i,t}$. If non-seasonal workers are in steady-state, with no changes in earnings or income on average, then the estimated change in each component will be equal to the change for seasonal workers attenuated by the fraction of seasonal workers, p. Importantly, even though these estimates

[33] Del Bono and Weber (2008) find that 64% of seasonal separations are ultimately followed by recall to the previous employer. Nekoei and Weber (2015) report a related exercise, also in Austria. Distinguishing between temporary and permanent layoffs on the basis of promised rehiring, they find that 23% of new employment contracts separated by a temporary layoff are spaced 12 months apart, compared with 13% of contracts separated by a permanent layoff.

of the changes are attenuated, our estimate of the pass-through rate will recover the true pass-through rate without attenuation:

$$\frac{\mathbb{E}\left[\Delta\text{Income}_{i,t} \mid \hat{\delta}_{i,t}=1\right]}{\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \hat{\delta}_{i,t}=1\right]} = \frac{p\mathbb{E}\left[\Delta\text{Income}_{i,t} \mid \delta_{i,t}=1\right]}{p\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \delta_{i,t}=1\right]} = \frac{\mathbb{E}\left[\Delta\text{Income}_{i,t} \mid \delta_{i,t}=1\right]}{\mathbb{E}\left[\Delta\text{Earnings}_{i,t} \mid \delta_{i,t}=1\right]} \tag{13}$$

The pass-through rate is still informative even if non-seasonal workers are not in steady-state. We can express the expected changes in earnings and household income among observations we classify as seasonal work interruptions—that is, those for which our proxy $\hat{\delta}$ equates to 1—as mixtures of these changes for seasonal and non-seasonal workers (suppressing subscripts for simplicity):

$$\begin{aligned} \mathbb{E}\left[\Delta\text{Earnings} \mid \hat{\delta}=1\right] &= p\mathbb{E}\left[\Delta\text{Earnings} \mid \delta=1\right] + (1-p)\mathbb{E}\left[\Delta\text{Earnings} \mid \delta=0\right] \\ \mathbb{E}\left[\Delta\text{Income} \mid \hat{\delta}=1\right] &= p\mathbb{E}\left[\Delta\text{Income} \mid \delta=1\right] + (1-p)\mathbb{E}\left[\Delta\text{Income} \mid \delta=0\right] \end{aligned} \tag{14}$$

where p is the fraction of observations with true seasonal work interruptions.

From these, we can define the measured pass-through rate (denoted θ^M) that we obtain by using our imperfect proxy, as well as the true pass-through rates for seasonal workers (θ^S) and non-seasonal workers (θ^{NS}):

$$\begin{aligned} \theta^M &\equiv \frac{\mathbb{E}[\Delta\text{Income} \mid \hat{\delta}=1]}{\mathbb{E}[\Delta\text{Earnings} \mid \hat{\delta}=1]} \\ \theta^S &\equiv \frac{\mathbb{E}[\Delta\text{Income} \mid \delta=1]}{\mathbb{E}[\Delta\text{Earnings} \mid \delta=1]} \\ \theta^{NS} &\equiv \frac{\mathbb{E}[\Delta\text{Income} \mid \delta=0]}{\mathbb{E}[\Delta\text{Earnings} \mid \delta=0]} \end{aligned} \tag{15}$$

Using these definitions, we can rewrite the measured pass-through rate that we estimate in the main text as the true rate of pass-through among seasonal workers plus a bias term:

$$\begin{aligned}
\theta^M &= \frac{\mathbb{E}[\Delta\text{Income} \mid \hat{\delta} = 1]}{\mathbb{E}[\Delta\text{Earnings} \mid \hat{\delta} = 1]} \\
&= \frac{p\mathbb{E}[\Delta\text{Income} \mid \delta = 1] + (1-p)\mathbb{E}[\Delta\text{Income} \mid \delta = 0]}{p\mathbb{E}[\Delta\text{Earnings} \mid \delta = 1] + (1-p)\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]} \\
&= \frac{p\theta^S\mathbb{E}[\Delta\text{Earnings} \mid \delta = 1] + (1-p)\theta^{NS}\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]}{p\mathbb{E}[\Delta\text{Earnings} \mid \delta = 1] + (1-p)\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]} \qquad (16) \\
&= \frac{p\theta^S\mathbb{E}[\Delta\text{Earnings} \mid \delta = 1] + (1-p)\theta^{S}\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]}{p\mathbb{E}[\Delta\text{Earnings} \mid \delta = 1] + (1-p)\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]} \\
&\quad + \frac{(1-p)\theta^{NS}\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0] - (1-p)\theta^{S}\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]}{p\mathbb{E}[\Delta\text{Earnings} \mid \delta = 1] + (1-p)\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]}
\end{aligned}$$

which simplifies to

$$\theta^M = \theta^S + \underbrace{\frac{(1-p)\left(\theta^{NS} - \theta^S\right)\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]}{\mathbb{E}[\Delta\text{Earnings} \mid \hat{\delta} = 1]}}_{\text{bias term}} \qquad (17)$$

The sign of the bias term depends on its components. Since $\hat{\delta}$ is a proxy for seasonal work interruptions, which lower own earnings, the denominator ($\mathbb{E}[\Delta\text{Earnings} \mid \hat{\delta} = 1]$) is negative. The numerator is positive if both $(\theta^{NS} - \theta^S)$ and ($\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]$) have the same sign, otherwise it is negative. Therefore, in order for θ^M to be a lower bound for θ^S (i.e., a negative bias term), $(\theta^{NS} - \theta^S)$ and ($\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]$) must be either both positive or both negative.

Since earnings tend to rise over a worker's career, $\mathbb{E}[\Delta\text{Earnings} \mid \delta = 0]$ is likely to be positive. Indeed, the evolution of earnings for workers least predicted to be seasonal shown in Figure 4 drifts upwards steadily throughout the 18 months following an initial separation.

Therefore, the sign of the bias term depends inversely on the sign of $\left(\theta^{NS} - \theta^S\right)$. If $\theta^{NS} > \theta^S$, as would be the case if seasonal workers are attempting to offset their changes in earnings via other margins within the household while non-seasonal workers are not, then the bias term will be negative and measured pass-through will be a lower bound for true pass-through among seasonal workers. The assumption that $\theta^{NS} > \theta^S$ is consistent with the results shown in Figure 4, where workers predicted to be non-seasonal have pass-through greater than 100% while workers predicted to be seasonal have pass-through less than 100%.

Appendix F Additional Results

Appendix Figure F.1: Regression-adjusted probabilities of recurrent separations

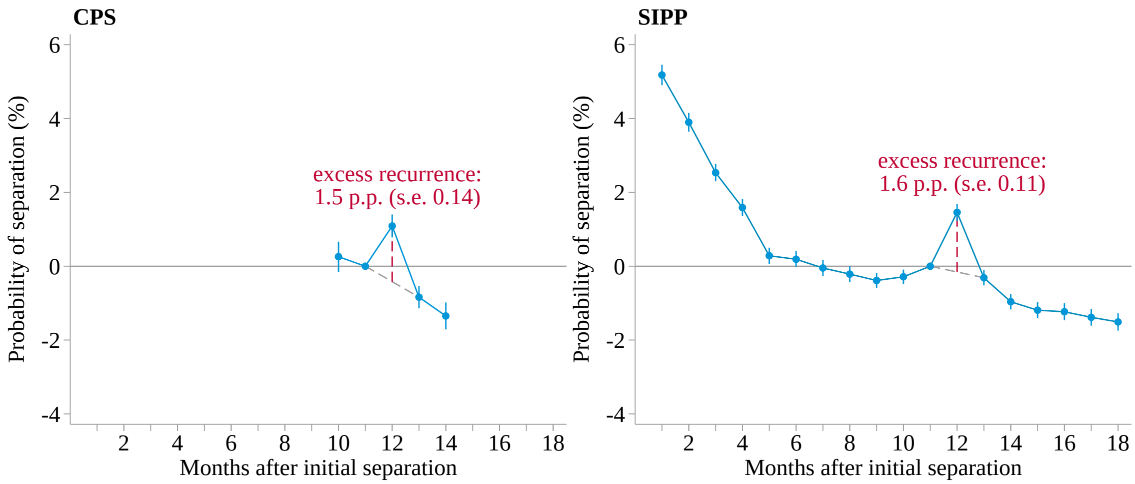

Notes: Probabilities $\hat{\rho}_\tau$ of experiencing a recurrent separation τ months later, estimated among prime-age CPS and SIPP workers observed exiting from employment in a baseline period. The CPS regression controls for the number of weeks between consecutive months' reference weeks. The SIPP regression controls for five-week months and for months coinciding with interview seams. We normalize the probability of separation to zero in month $\tau = 11$. Spikes show 95% confidence intervals, clustered by individual.

Appendix Figure F.2: Excess annual recurrence of job separations: by month of separation

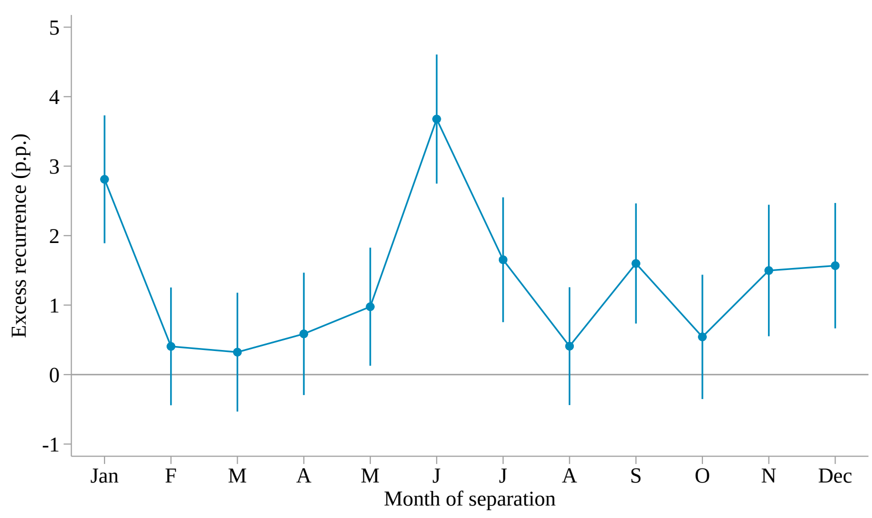

Notes: Excess annual recurrence of separations in our CPS sample, obtained by estimating our preferred specification separately by the month in which the original separation occurred. Spikes show 95% confidence intervals, clustered by individual.

Appendix Figure F.3: Excess annual recurrence of job separations: by pre-separation sector

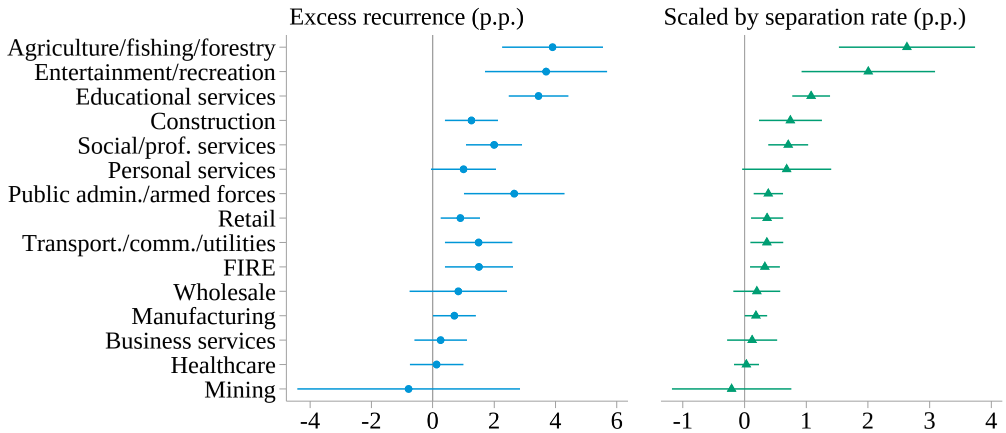

Notes: The left panel reports the excess annual recurrence of separations in our CPS sample, obtained by estimating our preferred specification separately by the industry a worker held prior to the initial separation. The right panel scales each estimate by 12 × that industry's monthly separation hazard. Both panels rank industries by this latter, scaled statistic. Spikes show 95% confidence intervals, clustered by individual.

Appendix Figure F.4: Excess recurrence of job separations among newly unemployed workers

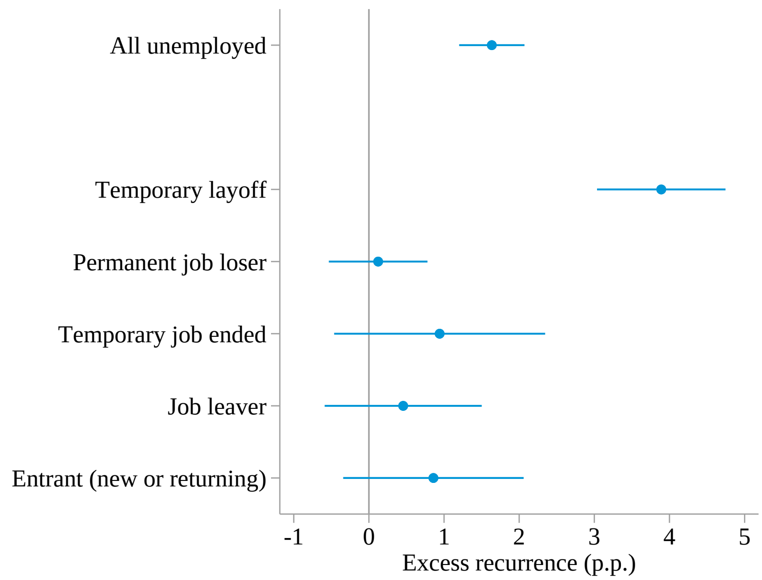

Notes: Estimated excess recurrence of job separations at a 12-month frequency among CPS respondents whose initial separation was into unemployment, and then subdividing these respondents by the reported reason for unemployment. In each case, we code all subsequent employment-to-non-employment transitions as recurrent separations, without regard for the type or reason for the second separation. Spikes show 95% confidence intervals, clustered by individual.

Appendix Figure F.5: Decomposition of recurrent separations using employer and industry identifiers

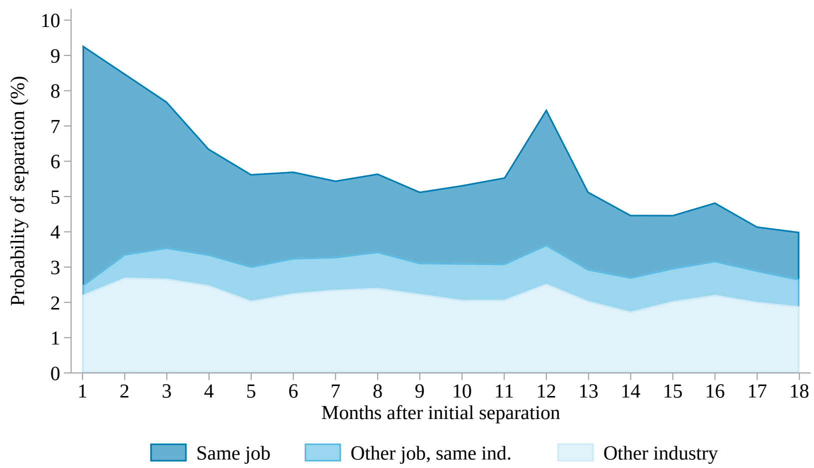

Notes: The sample consists of prime-age SIPP separators surveyed in the 1990–1993 SIPP panels. Using reliable employer identifiers available for these particular panels (Stinson, 2003), the figure additively decomposes the (unadjusted) probabilities $\hat{\rho}_\tau$ that a worker experiences a recurrent job separation τ months after the base separation into repeat separations (i) from the original employer, (ii) from a different employer in the original industry, or (iii) from a different employer in a different industry. Where either the base employer ID or the recurrent ID is missing, we assume that the two separations are made from different employers; where industry IDs are missing, we assume that recurrent separations are made from different industries.

Appendix Figure F.6: Predicted excess recurrence: distribution and validation

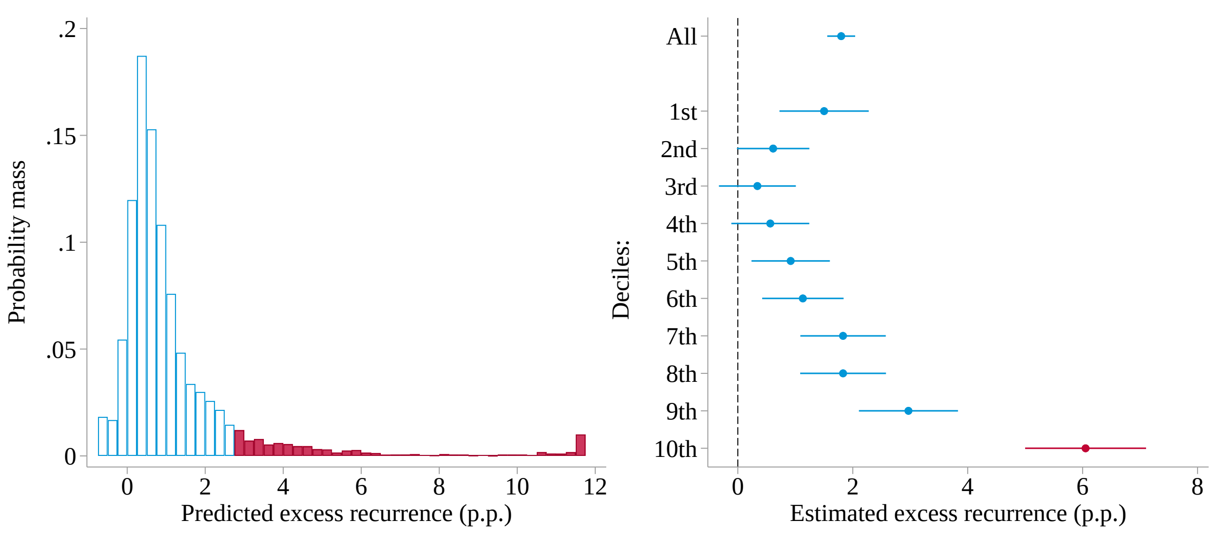

Notes: The left panel plots the distribution of predicted excess recurrence for separators in our SIPP sample, with the top decile shown in red. The right panel reports the estimated excess recurrence in the SIPP itself for each decile of this distribution. All estimates use our preferred specification. Spikes show 95% confidence intervals, clustered by individual.

Appendix Figure F.7: Earnings and household income of most/least-seasonal separators (denominated by household income at baseline)

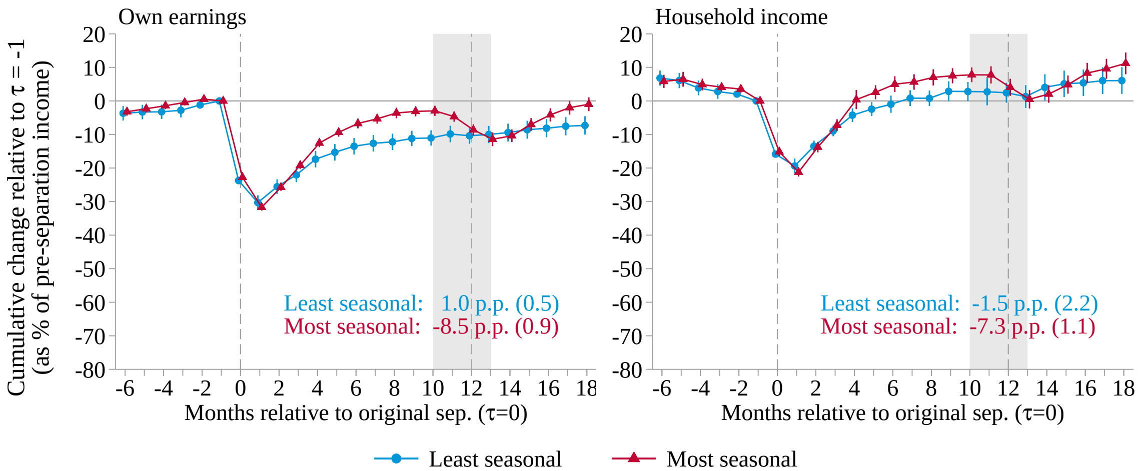

Notes: Left panel plots the evolution of personal earnings for prime-age SIPP respondents who experience an initial separation and are in either the top decile ("most seasonal") or bottom half ("least seasonal") of predicted excess recurrence. Right panel plots household income for this same sample. Both variables are expressed as a percentage of household income one month prior to the initial separation. Spikes show 95% confidence intervals, clustered by individual.

Appendix Figure F.8: Distribution of time to reemployment after the initial separation

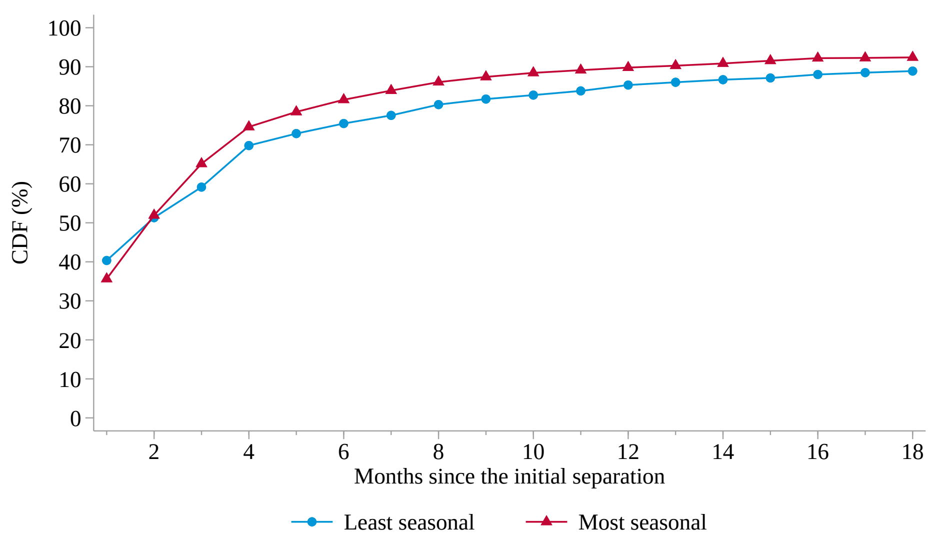

Notes: Cumulative distribution function of time to reemployment among prime-age SIPP separators who are in either the top decile ("most seasonal") or bottom half ("least seasonal") of predicted excess recurrence. Separators are coded as reemployed as soon as they have worked for at least one week since the initial separation.

Appendix Figure F.9: Employment status of most/least-seasonal separators

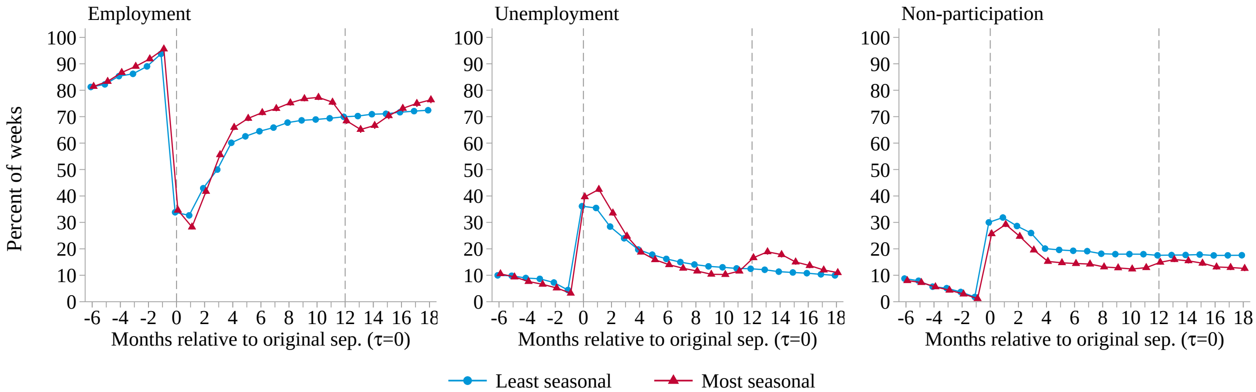

Notes: Each panel plots the average share of weeks spent in each employment status for prime-age SIPP respondents who experience an initial separation and are in either the top decile ("most seasonal") or bottom half ("least seasonal") of predicted excess recurrence.

Appendix Figure F.10: Total off-season deviations in household income (relative to pre-separation earnings)

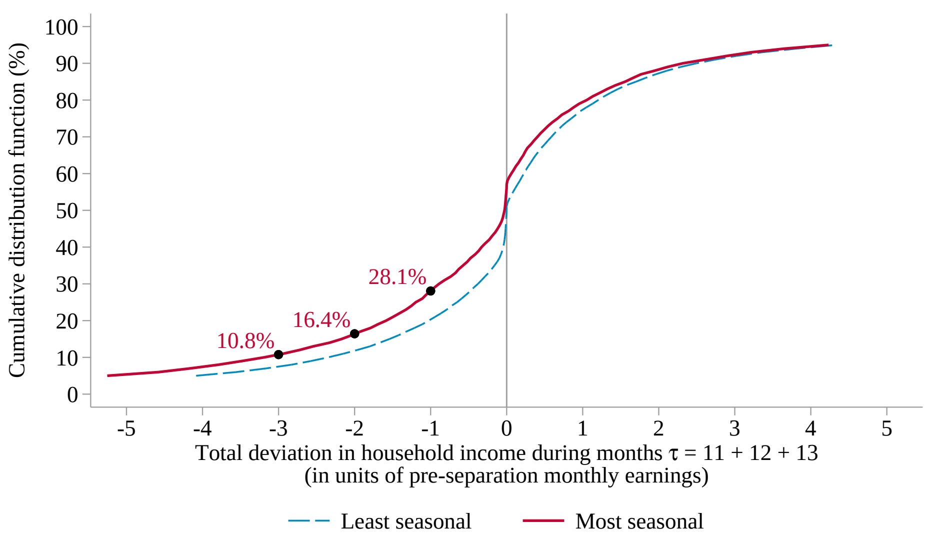

Notes: The sample consists of prime-age SIPP separators who are in either the top decile ("most seasonal") or bottom half ("least seasonal") of predicted excess recurrence. For each household i, we calculate the total off-season deviation in household income as $\Delta Y_i \equiv \sum_{\tau=11}^{13} \frac{y_{i,t_0+\tau} - y_{i,t_0+10}}{\text{Earnings}_{i,t_0-1}}$, where τ indexes months since the baseline job separation, then plot the empirical cumulative distribution function of ΔY_i across households.

Appendix Figure F.11: Earnings of seasonal separators in states with low vs. high unemployment rates

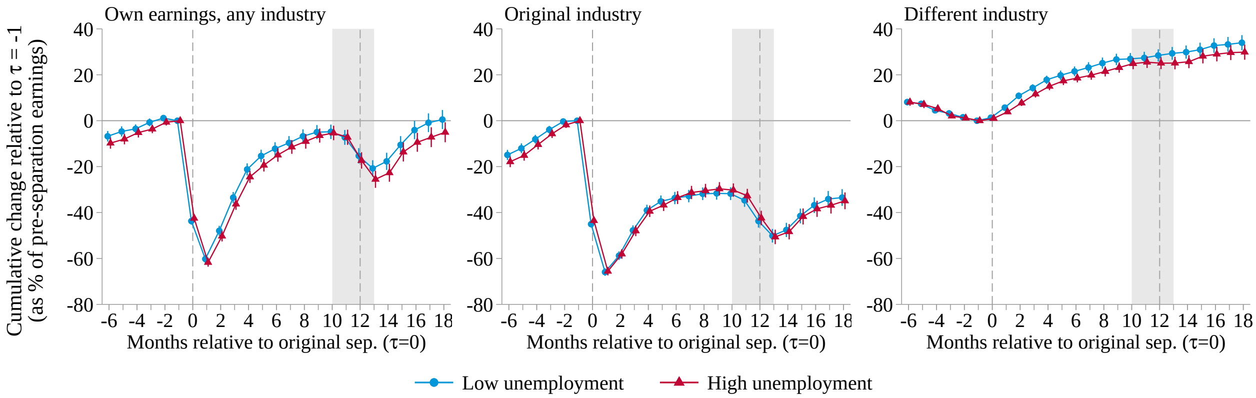

Notes: The left panel plots the evolution of personal earnings for prime-age SIPP respondents for whom the state unemployment rate at the time of the initial separation is below or above the sample medium. The middle and right panels decompose personal earnings based on the industry of the initial separation. Changes in earnings are expressed as a percentage of the individual's earnings one month prior to the initial separation. Spikes show 95% confidence intervals, clustered by individual.

Appendix Table F.1: Summary statistics: Current Population Survey (CPS) sample

	Employed workers (N = 15,383,408)		All job separators (N = 300,584)		Recurrent separators (N = 7,096)	
Demographics						
Female	45.9	(49.8)	55.5	(49.7)	57.3	(49.5)
Age	39.0	(8.4)	38.0	(8.5)	39.8	(8.2)
Non-white	26.5	(44.1)	36.5	(48.1)	33.3	(47.1)
Non-college	42.0	(49.4)	55.1	(49.7)	55.3	(49.7)
Household structure						
Spouse or partner in household	69.3	(46.1)	63.6	(48.1)	70.3	(45.7)
Spouse/partner employed (if present)	79.4	(40.5)	79.5	(40.4)	82.0	(38.4)
Spouse/partner in same industry (if employed)	22.2	(41.5)	22.0	(41.4)	21.5	(41.1)
≥1 children in household	54.2	(49.8)	57.3	(49.5)	61.2	(48.7)
Select industry indicators			*(For separators: measured pre-separation)*			
Agriculture, fishing, & forestry	2.3	(15.1)	4.6	(20.9)	8.3	(27.6)
Construction	7.2	(25.9)	12.4	(32.9)	14.9	(35.6)
Educational services	8.6	(28.1)	8.0	(27.1)	12.2	(32.7)
Healthcare	9.7	(29.6)	6.6	(24.8)	4.3	(20.3)
Separated into …						
Separated into unemployment			42.1	(49.4)	43.6	(49.6)
Separated into non-participation			57.9	(49.4)	56.4	(49.6)

Notes: All columns restrict to workers ages 25–54 observed during 1984–2013. “Employed workers”: restrict to person-months with employment during the reference week. “All job separators”: restrict to person-months for which the person is currently non-employed but was employed in the previous month. “Recurrent separators”: restrict to the subset of separator-months preceded by a similarly defined separation exactly 12 months prior. Separators’ industries are measured one month prior to separation. Except for age, all statistics are expressed as percentages (s.d. in parentheses).

Appendix Table F.2: Summary statistics: Survey of Income and Program Participation (SIPP) sample

	Employed workers (N = 8,085,096)		All job separators (N = 163,176)		Recurrent separators (N = 6,968)	
Demographics						
Female	46.2	(49.9)	52.0	(50.0)	48.9	(50.0)
Age	38.7	(8.3)	37.1	(8.3)	38.5	(8.0)
Non-white	24.0	(42.7)	28.5	(45.2)	28.6	(45.2)
Non-college	40.1	(49.0)	52.1	(50.0)	56.6	(49.6)
Household structure						
Spouse or partner in household	67.5	(46.8)	61.2	(48.7)	64.4	(47.9)
Spouse/partner employed (if present)	80.0	(40.0)	79.2	(40.6)	77.7	(41.6)
Spouse/partner in same industry (if employed)	20.6	(40.4)	18.0	(38.5)	19.7	(39.8)
≥1 children in household	54.5	(49.8)	57.0	(49.5)	58.8	(49.2)
Select industry indicators			*(For separators: measured pre-separation)*			
Agriculture, fishing, & forestry	1.3	(11.5)	3.0	(17.0)	6.3	(24.3)
Construction	5.4	(22.6)	11.1	(31.4)	16.5	(37.1)
Educational services	10.0	(30.0)	8.5	(27.9)	13.9	(34.6)
Healthcare	10.1	(30.2)	7.0	(25.5)	3.4	(18.1)
Monthly receipts (2017$)			*(For separators: measured pre-separation)*			
Personal earnings	3,750.9	(3,664.3)	2,019.7	(2,610.8)	1,749.5	(2,010.6)
Household income	6,970.4	(5,438.7)	5,254.6	(4,780.9)	4,897.2	(4,182.1)
Separated into …						
Separated into unemployment			56.2	(49.6)	66.1	(47.3)
Separated into non-participation			43.8	(49.6)	33.9	(47.3)

Notes: All columns restrict to workers ages 25–54 observed during 1984–2013. "Employed workers": restrict to person-months with continuous employment throughout the month. "All job separators": restrict to person-months in which we observe a week of non-employment immediately preceded (potentially in the prior month) by a week of employment. "Recurrent separators": restrict to the subset of separator-months preceded by a similarly defined separation exactly 12 months prior. Separators' industries, earnings, and income are measured one month prior to separation. With the exceptions of age, earnings, and income, all statistics are expressed as percentages (s.d. in parentheses).

Appendix Table F.3: Robustness of excess recurrence estimates

	Estimate	s.e.
CPS: no controls	1.44	0.14
SIPP: no controls	2.04	0.11
CPS: add regression controls	1.51	0.14
SIPP: add regression controls	1.62	0.11
SIPP: restrict to a balanced panel	1.54	0.13
SIPP: pre-1996 panels	1.68	0.16
SIPP: post-1996 panels	1.54	0.13
SIPP: excess recurrence of job-finding	1.50	0.10

Notes: Each row reports an estimate of the excess recurrence of job separations (or of job-finding) at annual frequencies. See Appendix C for details.

Appendix Table F.4: Properties of SIPP separations in highly seasonal industries

			Share of separations…		
Industry	Occupation	Share of industry	Classified as seasonal	Occurring in winter	Occurring in summer
Agriculture/fishing/forestry	Farming, forestry, and fishing	82.2	46.0	33.9	21.6
	Transportation and material moving	3.8	68.8	47.6	18.7
	Administrative support, including clerical	3.0	7.2	29.8	24.6
	Precision production, craft, and repair	1.7	28.0	40.2	21.8
Construction	Precision production, craft, and repair	69.0	8.7	35.0	20.3
	Handlers, equipment cleaners, helpers, and laborers	12.3	18.6	35.3	21.4
	Transportation and material moving	6.9	45.7	49.6	14.8
	Executive, administrative, and managerial	3.4	14.0	27.0	24.0
Entertainment/recreation	Services, except household and protective	49.8	21.3	25.1	24.9
	Professional specialty	13.6	16.1	25.0	28.5
	Executive, administrative, and managerial	11.5	12.0	23.3	28.3
	Sales	6.6	88.0	27.9	31.4
Educational services	Professional specialty	57.5	37.4	17.8	49.3
	Administrative support, including clerical	17.1	36.6	16.5	51.9
	Services, except household and protective	12.9	29.7	16.6	50.1
	Executive, administrative, and managerial	4.7	30.2	20.9	39.9

Notes: For each industry, we list the four occupations comprising the largest share of separations, followed by (i) the occupation's share of all separations from that industry, (ii) the share of separations in that industry–occupation cell that fall within the top decile of predicted excess recurrence, (iii) the share of separations occurring in early winter (November through January), and (iv) the share of separations occurring in early summer (May through July).

Appendix Table F.5: Changes in co-resident earnings by exposure to seasonal work interruptions

	All households		**Conditional on presence**	
	Most seasonal	Least seasonal	Most seasonal	Least seasonal
Partner earnings	-2.60	1.00	-4.68	0.70
	(0.97)	(0.39)	(1.38)	(0.60)
... in separator's industry	-2.29	0.14	-3.25	0.19
	(0.48)	(0.15)	(0.66)	(0.23)
... in different industries	-0.31	0.86	-1.43	0.51
	(0.90)	(0.38)	(1.29)	(0.58)
Other co-resident earnings	-1.13	-1.06	-6.13	-3.75
	(0.81)	(0.34)	(2.22)	(0.93)

Notes: Changes in household member earnings between 10 and 13 months after an initial separation for SIPP separators in the top decile ("most seasonal") and bottom half ("least seasonal") of predicted excess recurrence. The columns labeled "Conditional on presence" use only the observations with a spouse or unmarried partner (first three rows) or a non-partner household adult (last row). All earnings variables are expressed as percentages of the separator's earnings one month prior to the initial separation. Standard errors shown in parentheses.

Appendix Table F.6: Changes in transfers by exposure to seasonal work interruptions

	Most seasonal	Least seasonal
Household UI receipts		
Unadjusted	4.65	-0.28
	(0.40)	(0.12)
Adjusted for underreporting	6.93	-0.43
	(0.60)	(0.17)
Other transfer receipts	0.19	-0.04
	(0.13)	(0.05)
SNAP receipt (p.p.)	0.56	0.14
	(0.30)	(0.14)

Notes: Changes in government transfers between 10 and 13 months after an initial separation for SIPP respondents in the top decile ("most seasonal") and bottom half ("least seasonal") of the predicted excess recurrence distribution. Both transfer income variables are expressed as percentages of the separator's earnings one month prior to the initial separation. Food-stamp receipt is reported in percentage points. Standard errors shown in parentheses.

Appendix Table F.7: Decomposing changes in household income among the least seasonal separators

	Overall	Sex of focal separator: Male	Female	Type of separation: Unemp.	Non-part.	Pre-separation sector: Education	Other
A. Unadjusted estimates							
Own earnings	2.67	2.38	2.96	2.60	2.77	3.57	2.62
	(0.42)	(0.66)	(0.52)	(0.57)	(0.62)	(1.81)	(0.43)
Partner earnings	1.00	0.39	1.65	0.84	1.23	1.76	0.97
	(0.39)	(0.32)	(0.72)	(0.42)	(0.74)	(2.14)	(0.40)
Other co-resident earnings	-1.06	-0.78	-1.36	-1.15	-0.94	-3.93	-0.92
	(0.34)	(0.45)	(0.51)	(0.44)	(0.54)	(1.84)	(0.34)
Household UI benefits	-0.28	-0.26	-0.29	-0.35	-0.17	-0.42	-0.27
	(0.12)	(0.17)	(0.16)	(0.17)	(0.15)	(0.28)	(0.12)
Other transfer receipts	-0.04	-0.05	-0.03	-0.07	0.00	0.09	-0.05
	(0.05)	(0.06)	(0.08)	(0.06)	(0.09)	(0.13)	(0.05)
All other sources	-0.30	-0.23	-0.37	-0.11	-0.59	0.21	-0.32
	(0.23)	(0.29)	(0.35)	(0.26)	(0.40)	(1.12)	(0.23)
Household income	1.99	1.44	2.56	1.76	2.31	1.28	2.02
	(0.66)	(0.86)	(1.00)	(0.82)	(1.09)	(3.56)	(0.67)
B. Adj. for UI underreporting							
Household UI benefits	-0.43	-0.41	-0.45	-0.54	-0.26	-0.68	-0.41
	(0.17)	(0.24)	(0.24)	(0.25)	(0.21)	(0.42)	(0.18)
Household income	1.84	1.30	2.40	1.57	2.22	1.03	1.88
	(0.66)	(0.86)	(1.00)	(0.82)	(1.09)	(3.56)	(0.67)

Notes: Changes in components of household income between 10 and 13 months after an initial separation for SIPP respondents in the bottom half of predicted excess recurrence. All variables are expressed as percentages of the individual's earnings one month prior to the initial separation. Standard errors shown in parentheses.