12-2-2020

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

John Coglianese

Federal Reserve Board of Governors, john.m.coglianese@frb.gov

Brendan M. Price

Federal Reserve Board of Governors, brendan.m.price@frb.gov

Upjohn Institute working paper ; 20-337

Follow this and additional works at: https://research.upjohn.org/up_workingpapers

Part of the [Labor Economics Commons](#)

Citation

Coglianese, John and Brendan M. Price. 2020. "Income in the Off-Season: Household Adaptation to Yearly Work Interruptions." Upjohn Institute Working Paper 20-337. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/wp20-337>

This title is brought to you by the Upjohn Institute. For more information, please contact repository@upjohn.org.

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

Upjohn Institute Working Paper 20-337

John Coglianesi
Federal Reserve Board of Governors
Email: john.m.coglianesi@frb.gov

Brendan M. Price
Federal Reserve Board of Governors
Email: brendan.m.price@frb.gov

November 2020

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 U.S. workers experience recurrent separations spaced exactly 12 months apart. These separations coincide with aggregate seasonal downturns and are concentrated in seasonally volatile industries. Examining workers most prone to seasonal work interruptions, we find that these workers incur large earnings losses during the off-season. Lost earnings are 1) driven mainly by repeated separations from the same employer, 2) not recouped at other firms, 3) partly offset by unemployment benefits, and 4) amplified by concurrent drops in partners' earnings. On net, household income falls by about \$0.80 for each \$1 lost in own earnings.

JEL Classification Codes: D10, E32, J63

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

Acknowledgments: 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 U.S. 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. All errors are ours. 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.

Upjohn Institute working papers are meant to stimulate discussion and criticism among the policy research community. Content and opinions are the sole responsibility of the author.

1 Introduction

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 in earned income at particular points in the calendar year. [Figure 1](#) offers one illustration of this fact: the unemployment rate among prime-age U.S. 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 percent 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.² Using three decades of panel data on U.S. workers, we find robust empirical evidence that an excessive share of workers undergo repeated transitions from employment to nonemployment spaced exactly 12 months apart. These recurrent transitions—which for brevity we call *separations*—appear to be seasonal in origin, as they align closely with temporal, sectoral, and demographic features of the seasonal cycle. To examine whether households recoup the income lost due to seasonal work interruptions, we use preseparation 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.

¹Own 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.

²Precursors to our approach appear in [de Raaf et al. \(2003\)](#) 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.

Figure 1: Nonseasonally adjusted unemployment among prime-age U.S. workers, 1984–2013

Notes: The left panel plots the raw monthly unemployment rate among U.S. 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 percent of workers can be matched to the prior month.

We begin by developing a simple procedure for detecting and analyzing seasonal work interruptions. Under the testable assumption that seasonally employed workers tend to lose their jobs at the same time each year, we show how high-frequency panel microdata can be used to identify seasonal separators on the basis of repeated separations spaced 12 months apart. Using the Survey of Income and Program Participation (SIPP), we find that prime-age workers are indeed disproportionately likely to exit employment exactly 12 months after an initial separation. Our preferred estimates indicate that separations are 1.6 percentage points more likely to be repeated at 12-month intervals than would be expected on the basis of separation rates 11 or 13 months after an initial separation. Our baseline estimate of the excess mass of separations that repeat at yearly intervals—which we term *excess recurrence* for short—is robust to a host of sensitivity and validation exercises. For example, we replicate our excess recurrence calculation in the Current Population Survey (CPS) and obtain a similar estimate of 1.4 percentage points. Moreover, among separations into unemployment, excess recurrence is driven by CPS respondents who report being job losers, rather than by self-described job leavers.

Though nothing in our method draws a mechanical link to the aggregate seasonal cycle, the

yearly work interruptions isolated by our method correspond closely with seasonal patterns in the broader labor market. First, consistent with the idea that seasonal job destruction triggers worker dislocation, excess recurrence is strongest in the calendar months in which aggregate employment tends to contract. Second, excess recurrence is concentrated in industries that are subject to seasonal shifts in labor demand, such as agriculture, education, recreation, and construction. Third, excess recurrence is most pronounced for the demographic groups whose employment rates fluctuate most widely over the calendar year. We find evidence of a negative skill gradient in exposure to seasonal work interruptions: with the exception of teachers (who are often college-educated), repeated yearly separations occur more often among less-educated workers.

Having validated that annually recurrent separations are indicative of seasonal work interruptions, we next ask to what extent households exposed to seasonal earnings losses experience net reductions in income as a whole. We begin by charting the evolution of personal earnings and household income among workers who experience two separations into nonemployment spaced exactly 12 months apart. For this initial sample, we find steep declines in earnings that largely pass through to lower income: for every \$1 drop in a worker’s own earnings at the time of the second separation, her household income drops by \$0.78, suggesting that other forms of income-generation only modestly offset the decline in income induced by a seasonal job loss.

However, examining earnings and income dynamics among workers with *observed* separations into nonemployment recurring 12 months apart may yield misleading inferences about household adaptation to seasonal work interruptions. Workers who anticipate being subject to seasonal layoffs may in some cases secure new jobs to begin as soon as their seasonal jobs end. Excluding such workers from consideration potentially obscures the degree to which individuals subject to seasonal work interruptions stabilize their earnings by taking on alternative employment during the off-season. Conditioning on an observed period of joblessness could also yield biased estimates of the extent to which seasonal reductions in earnings pass through to household income.

To circumvent this concern, we therefore construct an alternative sample of workers who are *prone* to seasonal work interruptions, without conditioning on whether a repeat exit from employment is actually realized. To do so, we first nonparametrically estimate variation in excess recurrence among CPS separators on the basis of their preseparation industry and occupation, the time of year, and local temperature. We then use these estimates to predict the excess likelihood

of an annually recurrent separation for each separator in our primary SIPP sample. This approach is able to capture subtle patterns of seasonal work across job types, such as the pattern that construction workers who exit from employment in December are very likely to separate again the following December if they reside in a cold state, but not especially likely if they reside in a warm state. We assign SIPP separators to deciles based on their predicted excess recurrence, then compare earnings and income dynamics among separators *most* likely (top decile) to experience an annual recurrence to those among separators who are *least* likely (bottom decile) to do so.

This more-representative analysis confirms that, on net, households recoup only a small share of the earnings lost in the course of seasonal work interruptions. Among job-separators most likely to experience seasonal interruptions, personal earnings fall by an amount equal to 18.6 percent of preseparation earnings between the 10th and 13th months following their initial separation, as a significant fraction of these workers undergo a recurrent nonemployment spell. During this same period, these workers' household incomes fall by an amount equal to 15.1 percent of preseparation earnings, implying a pass-through rate of \$0.81 for each \$1 in foregone earnings.³ Our finding that the large majority of lost earnings are passed through to household income is broadly shared across different groups of seasonal separators, including men and women, educators and noneducators, those who initially separated into unemployment, and those who separated into nonparticipation. By contrast, workers with low estimated seasonal exposure continue to recover earnings and household income between 10 and 13 months after their original separation.

To unpack the large but incomplete earnings-to-income pass-through we observe in our data, we decompose the response of household income along several margins and uncover offsetting shifts in both directions. Declines in separators' earnings during *recurrent* separations are primarily driven by reductions in earnings at the job held before a worker's *initial* separation—a fact consistent with many recalls from temporary layoff in the population we study. These earnings losses are augmented by additional losses at other employers in the separator's original industry, and they are not offset by earnings gains in other sectors. Next, we find no evidence that spouses or other household members increase their labor supply during the off-season: surprisingly, we find modest *declines* in the earnings of spouses and unmarried partners that augment the overall drop in income. These

³These estimates are adjusted for underreporting of unemployment insurance receipts, using reporting rates provided by Meyer et al. (2015). Without adjustments, household income falls by 17.3 percent of preseparation earnings, passing through \$0.93 for each \$1 of lost earnings.

declines are driven by earnings in the same industry as the seasonal separator, indicating that many seasonal households experience correlated earnings fluctuations between partners. Turning lastly to transfers, we find that unemployment insurance replaces about one-third of the earnings lost by seasonal separators. We also find a modest, marginally significant increase in receipt of Supplemental Nutrition Assistance Program benefits. Means-tested cash transfers, by contrast, do not appear to play a significant role in insulating households against seasonal income losses.

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 et al., 1993](#); [Couch and Placzek, 2010](#); [Davis and von Wachter, 2011](#); [Lachowska et al., 2020](#)). A separate series of papers has catalogued the prevalence of temporary layoffs followed by workers’ eventually being recalled to their former employers (e.g., [Feldstein, 1975](#); [Katz, 1986](#); [Katz and Meyer, 1990](#); [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 U.S. 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 et al., 2008](#); [Sabelhaus and Song, 2009](#); [Kopczuk et al., 2010](#); [DeBacker et al., 2013](#)). Recent work by [Farrell and Greig \(2015\)](#), [Hannagan and Morduch \(2015\)](#), and [Morris et al. \(2015\)](#) examines within-year income volatility experienced by households but does not separate seasonal variation from nonseasonal 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](#); [Spletzer, 1997](#); [Stephens, 2002](#); [Juhn and Potter, 2007](#)) and from government transfers (e.g., [Gruber, 1997](#); [Browning and Crossley, 2001](#); [Ganong and Noel, 2019](#)). In contrast to prior research suggesting that an “added worker effect” mitigates the drop in household income

incurred by displaced workers 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 the macroeconomic causes and consequences of seasonality. In a seminal paper, Barsky and Miron (1989) showed that economic activity is highly seasonal across a variety of different indicators, and that these seasonal fluctuations bear important resemblance to business cycle fluctuations. Subsequent work has clarified the relationship between the seasonal and business cycles (Beaulieu et al., 1992; Beaulieu and Miron, 1992; Miron and Beaulieu, 1996; Geremew and Gourio, 2018) and explored the pronounced seasonal patterns evident in certain segments of the economy, such as retail (Warner and Barsky, 1995), housing (Ngai and Tenreyro, 2014), and inventories (Miron and Zeldes, 1988). Relative to this literature, we examine the implications of seasonality for individual 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).

Our estimates imply that households do not engage in appreciable “income smoothing” in connection with seasonal work interruptions (Morduch, 1995). Of course, households may still smooth consumption either by borrowing or by drawing down savings. 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, 2019). 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 et al., 2014).

The rest of the paper proceeds as follows. Section 2 explains our approach of identifying seasonal work interruptions on the basis of annually recurrent separations. Section 3 describes our data. Section 4 establishes the empirical fact that an excess number of separations recur at

12-month intervals. [Section 5](#) explores heterogeneity in this phenomenon across different groups of workers and types of separations. [Section 6](#) examines the evolution of personal earnings and household incomes before and after seasonal work interruptions. [Section 7](#) tracks the margins along which households do or do not adapt to seasonal separations. [Section 8](#) concludes.

2 Measuring Seasonal Work

How can we identify seasonal work? If commonly used household surveys were to ask respondents to classify their jobs as seasonal or nonseasonal, then this question might be straightforward to answer, but unfortunately this is typically not the case. Even if such information were available, however, it is not clear how workers' self-reported seasonal status would correspond to economists' notions of seasonality. For example, a teacher who works in a second job over the summer break might report that neither job is seasonal, since from their perspective they are employed year-round, but both jobs are examples of seasonal work that we aim to examine in this paper.

Another possible approach would be to start from month-to-month changes in aggregate employment, perhaps focusing on sectors that exhibit large swings in employment synchronized with the economy-wide seasonal cycle. To identify likely seasonal workers, one could compare separations occurring at the onset of the low season (say, in January) to those occurring at other points of the year. However, defining seasonality in relation to aggregate patterns may obscure the prevalence of seasonality at the level of individual workers. For example, a construction laborer might work from June to August while a school bus driver works from September to May. Put together, their combined employment is flat throughout the year, even though both hold seasonal jobs. More generally, to the extent that different seasonal workers are employed during different portions of the calendar year, the seasonality of aggregate employment may be a poor guide to the extent of seasonal work as experienced by individuals. Classifying separations based on the timing of the aggregate seasonal cycle could lead one to ignore many seasonal workers whose jobs simply don't line up with this cycle.

Instead of either of these approaches, we identify seasonality on the basis of periodicity in individual work histories at annual frequencies. Our method is predicated on the idea that seasonal workers will tend to separate from jobs at the same point each year, while nonseasonal workers will

not. Specifically, we take a sample of individuals who separate from employment at some point in time and calculate their likelihood of experiencing recurrent separations at various points in the future. Seasonal workers should be excessively likely to experience another separation exactly 12 months later, while nonseasonal separators will be no more likely to experience a separation at an annual horizon than at other, similar spans of time. Therefore, the *excess recurrence* of separations at 12-month spans provides an estimate of seasonal work as experienced by individuals.

Given longitudinal data, we can estimate our measure of seasonal work using a simple regression approach. For an individual i who separates from employment into nonemployment in time period t_0 , let $y_{i,t} \equiv \mathbb{1}\{\text{separate in month } t\}$ be a binary variable denoting whether she experiences another such separation in period t . We estimate regressions of the form

$$y_{i,t} = \sum_{\tau} \rho_{\tau} \mathbb{1}\{t - t_0 = \tau\} + \mathbf{x}'_{i,t} \beta + \varepsilon_{i,t} \quad (1)$$

where $\mathbf{x}_{i,t}$ is a vector of control variables. The coefficient ρ_{τ} estimates the likelihood, conditional on the controls, of experiencing an additional separation exactly τ months after an initial separation. By estimating ρ_{τ} over a range of $\tau \in \{1, \dots, \tau_{\max}\}$, we track the evolution of this rate of recurrence following a baseline separation. We compute the excess recurrence at a 12-month span as $\rho_{12} - \frac{1}{2}(\rho_{11} + \rho_{13})$, which furnishes a measure of seasonality as experienced by individuals. This approach can be implemented in any panel dataset that tracks individual workers' employment statuses at high frequency for periods in excess of one year. It can also be used to compare seasonality in different segments of the labor market by estimating excess recurrence among different subsamples.

To formalize the intuition undergirding our approach, [Appendix A](#) describes a stylized economy in which a mixture of seasonal and nonseasonal workers cycle in and out of work. Importantly, this model clarifies that the excess recurrence of separations at annual frequencies identifies the fraction of *separations* that are seasonal, rather than the fraction of *workers*. When comparing different segments of the labor market subject to differing degrees of churn (as measured by separations per worker), it is helpful to rescale estimates of excess recurrence by the separation rate in each segment so that each estimate is on the same scale. We employ this adjustment in later sections when we examine heterogeneity in excess recurrence.

3 Household Panel Data

To identify which households are exposed to seasonal work interruptions—and 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. The Survey of Income and Program Participation (SIPP) affords us a comprehensive, longitudinal accounting of households’ work activities, earnings, and nonlabor income over time periods straddling multiple seasonal cycles. We complement the SIPP using the CPS, which—though less informative about household adjustment—furnishes larger samples that are well-suited to detecting annually recurrent separations. [Appendix B](#) offers additional detail about how we construct each sample.

3.1 Primary sample: the Survey of Income and Program Participation (SIPP)

Launched in 1984, the SIPP is a household survey produced by the U.S. 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 chart postseparation employment, earnings, and income dynamics for every month in a year-and-a-half period following an initial exit from employment.

Using SIPP 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.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 with our (monthly) income data, we first code a job separation as occurring in some week w if an individual reported

⁴While we confine our attention to seasonal work interruptions observed among prime-age workers, we retain co-residents of all ages when examining concurrent changes in income earned by other members of the household.

being employed in week $w - 1$ but nonemployed 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 all of our SIPP-based analyses, our notion of separations thus encompasses all new jobless spells that last for at least one week. Though we often speak of “job separations” for concreteness, our definition excludes job-to-job transitions that are unaccompanied by an intervening week of nonemployment.

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. Individuals may be less likely to remember the exact timing of employment changes further in the past, leading to higher rates of recorded labor force transitions between interview sessions than within them (Kalton et al., 1998). This phenomenon—commonly referred to 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. We describe in Section 4 how we correct for seam bias in our estimates.

Table 1 reports summary statistics for our SIPP sample, separately for continuously employed workers, workers undergoing a job separation, and the subset of separators who experience a second job separation exactly 12 months after the first one. Relative to the typical worker, separators are less likely to be college-educated, and they have below-average earnings and household incomes in the month prior to separation. These patterns are even stronger among annually recurrent separators, hinting that seasonal workers may be negatively selected on baseline earnings potential.⁷ Among job separations, 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 health care, a sector with comparatively stable employment levels throughout the calendar year.

⁵The 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 “nonparticipants” if they are *jobless and not looking for work*.

⁶For 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 she first returns to work in mid-March, then loses her job and exits the labor force in mid-April, she has a separation in April but no separation in February or March.

⁷Prior work has found that seasonal workers receive positive compensating wage differentials for the disamenity of seasonal job loss (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 Table 1 imply.

Table 1: Summary statistics: Survey of Income and Program Participation (SIPP) sample

	Employed workers		All job separators		Recurrent separators	
	(N = 8,085,096)		(N = 163,176)		(N = 6,968)	
<hr/>						
Demographics						
<hr/>						
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)
<hr/>						
Household structure						
<hr/>						
Married	66.6	(47.2)	60.2	(48.9)	62.7	(48.4)
≥1 children in household	54.5	(49.8)	57.0	(49.5)	58.8	(49.2)
<hr/>						
Select industry indicators			<i>(For separators: measured pre-separation)</i>			
<hr/>						
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)
<hr/>						
Monthly receipts (2017\$)			<i>(For separators: measured pre-separation)</i>			
<hr/>						
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)

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 nonemployment 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).

3.2 Auxiliary sample: the Current Population Survey (CPS)

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. We are thus able to estimate excess recurrence in the CPS as well as in the SIPP. Furthermore, because households are interviewed at monthly intervals, the CPS is free from the seam bias that arises in the SIPP.

While the CPS is well suited to pinpointing annually recurrent job separations, it is uninformative about how earnings and income fluctuate within households from month to month, as for

each household these outcomes are measured only once per year as part of the outgoing rotation group module. For this reason, we treat the SIPP as our primary dataset, relying on the larger CPS to corroborate our findings and to help predict which job separations are likely to recur 12 months later. As detailed in [Appendix B](#), we construct a monthly panel of prime-age CPS individuals over 1984–2013, matching our SIPP sample. We report CPS summary statistics in [Appendix Table D.1](#).

4 Annually Recurrent Work Interruptions

In this section, we establish empirically that individuals who separate from a job in a given month are indeed 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. We then demonstrate the close connections between annually recurrent job separations and characteristic features of labor market seasonality.

4.1 Excess recurrence of job separations at annual frequencies

We begin by using our SIPP sample to estimate the time path of recurrent job separations in the wake of an initial separation from employment into nonemployment. Recall from [Section 2](#) that ρ_τ denotes the probability that a worker who separates in period t_0 experiences a repeat separation τ months later, conditional on control variables. By estimating ρ_τ over a range of values of $\tau \in \{1, \dots, \tau_{\max}\}$ using [Equation 1](#), we track the evolution of this rate of recurrence in the aftermath of a baseline separation.⁸

The left panel of [Figure 2](#) presents our estimates of $\hat{\rho}_\tau$ for $\tau \in \{1, 2, \dots, 18\}$, tracing out the likelihood of a recurrent job separation over the year and a half following an initial separation.⁹ Job separations indeed appear to be autocorrelated at annual frequencies: the separation rate exactly 12 months after an initial separation is 7.6 percent, more than 1.5 percentage points (p.p.) greater than that at any nearby horizon. To measure the excess recurrence of job separations at an annual

⁸To 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.

⁹Note that, because our measure of job separations is based on underlying weekly detail, it is possible (and indeed fairly common) for SIPP respondents to register job separations in back-to-back months.

Figure 2: The probability of a recurrent job separation

Notes: Probabilities $\hat{\rho}_\tau$ of experiencing a recurrent separation τ months later, estimated in a sample of prime-age SIPP workers observed exiting from employment in a baseline period. The left panel presents the unadjusted probabilities. The right panel plots the coefficients from our preferred specification, which includes controls for five-week months and for months coinciding with interview seams, as described in the text. For ease of presentation, we normalize the coefficients so that $\frac{1}{2}(\hat{\rho}_{11} + \hat{\rho}_{13}) = 0$. Spikes show 95% confidence intervals, clustered by individual.

horizon, we 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 nonseasonal dynamics. Deducting the average of the 11th and 13th recurrence probabilities from the 12th gives us our first estimate of excess recurrence: $\hat{\rho}_{12} - \frac{1}{2}(\hat{\rho}_{11} + \hat{\rho}_{13}) = 2.0$ p.p.

However, some of the excess separations measured at $\tau = 12$ may be unrelated to seasonality. The observant reader will note that the separation probabilities for months 8, 12, and 16 are all elevated relative to those at neighboring horizons, a fact that we interpret as stemming from seam bias in the SIPP. As others have documented previously, 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). 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.¹⁰

To account for this possibility, our preferred specification—presented in the right panel of [Figure 2](#)—includes three control variables to address mechanical factors that may induce spurious autocorrelation at annual frequencies. First, to account for the possibility of seam bias, we include 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 for each respondent’s first monthly observation that coincides with a seam, and another indicator variable for all subsequent observations that coincide with seams. Second, to account for the fact that some calendar months are longer than others (with longer months affording more opportunities for a work interruption), we include an indicator variable for month-years that the SIPP regards as five weeks in length.¹¹ Incorporating these controls eliminates the excess separations observed at $\tau = 8$ and 16 , and it yields an adjusted estimate of excess recurrence equal to 1.6 p.p. This estimate is highly statistically significant, with a standard error of 0.11. All remaining SIPP-based estimates reported in this section are inclusive of these control variables.

A final feature of [Figure 2](#) warrants mention: we observe slightly elevated separation probabilities at $\tau = 11$ and $\tau = 13$ alongside the much larger spike at $\tau = 12$. This is to be expected: the exact timing of annually recurrent separations is likely to deviate slightly from 12 months for some workers due to slight year-to-year differences in the timing of workweeks and the onset of adverse winter weather, as well as imperfect recall of jobs’ exact start and end dates.¹² Insofar as the recurrence probabilities at months 11 and 13 are themselves elevated due to seasonality, a broader measure of seasonal work interruptions would incorporate the excess mass at horizons 11, 12, and 13 relative to horizons 10 and 14. Computing the statistic $\sum_{\tau \in \{11, 12, 13\}} (\hat{\rho}_{\tau} - \frac{1}{2}(\hat{\rho}_{10} + \hat{\rho}_{14}))$, we

¹⁰While [Figure 2](#) reveals no similarly elevated separation probability four months after baseline, examination of the underlying week-level data shows a less pronounced, but qualitatively similar spike occurring at the four-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).

¹¹The SIPP encodes each monthly observation as being four or five weeks long, so that there are either four or five weekly employment reports per person-month. Which calendar months last for five weeks varies from year to year based on the particular timing of days within the month, but 31-day months are more commonly coded as five-week months. This leads to a mechanical positive autocorrelation in length-of-month at 12-month frequencies, which imparts a slight upward bias to the raw recurrence probability at $\tau = 12$.

¹²Similar “spillovers” are notably *absent* next to the nonseasonal seams at $\tau \in \{4, 8, 16\}$, a useful check that the 12-month spike differs qualitatively (as well as quantitatively) from those at other multiples of four.

Table 2: Robustness of excess recurrence to alternative samples and specifications

	Estimate	s.e.
No controls	2.04	(0.11)
Preferred specification	1.62	(0.11)
Restrict to a balanced panel	1.54	(0.13)
Pre-1996 panels	1.68	(0.16)
Post-1996 panels	1.54	(0.13)
Replicate in the CPS	1.44	(0.14)
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. Except where indicated, all specifications are estimated using our SIPP estimation sample of prime-age workers. See text for details.

find a larger point estimate of 3.0 p.p. (with s.e. = 0.21), suggesting that our preferred estimate of excess recurrence is a conservative lower bound on the true prevalence of seasonal separations.

4.2 Robustness

Our excess recurrence measure survives a suite of sensitivity checks. Summarized in [Table 2](#), these checks confirm that the excess recurrence of job separations is a robust empirical fact.

Sample attrition To maximize statistical power, our preferred specification retains all available observations in the 18 months following each initial separation, even if an individual exits from the SIPP 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—a uniform source of Type I right-censoring that is unlikely to bias our estimates ([Kiefer, 1988](#))—or because of survey attrition, which is potentially nonrandom.

To assess whether dynamic changes in sample composition are biasing our estimates, we reestimate 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. The rate of excess recurrence in the later panels, although below that of the earlier panels, is still high and close to our overall estimate.

Replicating in the CPS We also check whether excess recurrence is apparent in our auxiliary CPS sample. The CPS records whether respondents reported being employed during each month’s reference week (typically the second week of the month). Accordingly, we define a job separation as a month of self-reported nonemployment preceded by a month in which the individual reports being employed. This separation concept differs slightly from the definition we use in the SIPP and will sometimes miss separations that are followed by brief jobless spells lasting less than one month. Nonetheless, it will capture the large majority of separations observed in the SIPP. Though the design of the CPS prevents us from estimating the profile of recurrent separations at all horizons, we can estimate a variant of [Equation 1](#) for $\tau \in \{10, 11, 12, 13, 14\}$, which is sufficient to compute excess recurrence. In congruence with our preferred SIPP specification, we control for the number of weeks elapsed between successive reference weeks, as more separations occur over longer intervals.¹³

Like SIPP separators, CPS separators experience excess recurrence of separations at 12-month spans. We estimate an excess recurrence rate of 1.4 p.p. (s.e. = 0.14) for our CPS sample, only slightly lower than our preferred estimates from the SIPP. The similarity of our SIPP- and CPS-based estimates is reassuring given the differences in measurement issues (e.g., seam bias) and employment concepts (i.e., weekly vs. reference week) between these two datasets.

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,

¹³The gap between reference weeks varies both month-to-month and year-to-year depending on how weeks and holidays are distributed within each month (see [Appendix B](#)). Though omitting this control could potentially bias our estimates, in practice we obtain a nearly identical estimate of excess recurrence in a regression without controls.

we compute the probability of a recurrent transition from nonemployment 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. Net of our preferred SIPP controls, 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.

4.3 Connecting recurrent separations to labor market seasonality

To what extent does the periodicity in *individual* job separations align with seasonal patterns at the *aggregate* level? To answer this question, we examine how the likelihood of a recurrent separation varies both by calendar month and by preseparation sector. Though nothing in our method creates a mechanical link between excess recurrence and the calendar year, we find close connections between these two phenomena.

By month of separation Net of secular trends, the employment-to-population rate among prime-age workers typically declines both at the beginning of winter and at the beginning of summer.¹⁴ 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.

Figure 3 uses our SIPP sample to estimate excess recurrence separately by the month in which the original separation occurred. Consistent with the early-summer and early-winter declines in aggregate employment, we find that individuals who experience job separations in June, November, and December are the most (excessively) likely to separate again 12 months later.^{15,16}

¹⁴We base this statement on a regression of the 1984–2013 prime-age CPS employment rate on month fixed effects, controlling for secular and business-cycle fluctuations using a restricted cubic spline in calendar time. Employment contracts 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.).

¹⁵Although CPS employment declines sharply in January, the SIPP shows little excess recurrence among January separations. This incongruity reflects a quirk of the CPS design: although the reference week normally straddles the 12th day of the month, the December reference week is routinely advanced by one week so that interviews conclude before Christmas. As a result, an outsized share of December separations first appear in the CPS in January. If we mimic the CPS separation concept in the SIPP by designating a pseudo-reference week for each month, we find less excess recurrence in December and more in January, reconciling the temporal patterns across datasets.

¹⁶Appendix Figure D.1 presents analogous estimates scaled by the average separation rate in each month. These scaled estimates convey the probability that a given worker will experience a seasonal separation in a given month.

Figure 3: Excess annual recurrence of job separations: by month of separation

Notes: Excess annual recurrence of job separations in our SIPP 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.

By preseparation industry The left panel of Figure 4 shows how excess recurrence varies across 15 one-digit sectors. Three of the leading sectors—1) agriculture, fishing, and forestry; 2) educational services; and 3) construction—line up neatly with our intuitions about seasonal rhythms in farming, schooling, and building.¹⁷ By contrast, finance, insurance, and real estate—in which employment counts fluctuate little throughout the year—shows no excess recurrence of individual separations. Most sectors exhibit some degree of excess recurrence, though, indicating that seasonality is somewhat widespread and not exclusively confined to a few small sectors.

The excess recurrence for a given industry tells us the *conditional* probability that a separation from that industry will be repeated one year later. Multiplying this statistic by the monthly separation rate within the industry yields an estimate of the *unconditional* probability that a given worker will experience a seasonally recurrent separation. The right panel of Figure 4 presents estimates of

$$\hat{\alpha}_i \equiv \text{excess recurrence for industry } i \cdot \widehat{\Pr}(\text{Sep} \mid \text{Industry} = i) \cdot 12 \quad (2)$$

¹⁷Interestingly, 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 wait around for the next holiday boom.

Figure 4: Excess annual recurrence of job separations: by pre-separation sector

Notes: The left panel reports the excess annual recurrence of job separations in our SIPP 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 $12 \times$ that industry's monthly separation hazard. Both panels rank industries by this latter, scaled statistic. Spikes show 95% confidence intervals, clustered by individual.

Scaling the 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.

Summary The consistency of these patterns with those of aggregate seasonality is reassuring. Our method for detecting annually recurrent job separations isolates work interruptions that coincide in time with seasonal downturns in aggregate employment and arise in sectors that are clearly prone to seasonal forces. For these reasons, we think that household responses to the seasonal work interruptions isolated by our method are likely to be informative about household adaptation to labor market seasonality more broadly.

5 Heterogeneity in Recurrent Separations

What kinds of job separations tend to recur at 12-month intervals? In this section, we explore heterogeneity along three dimensions. First, stratifying by education, we find that excess recurrence exhibits a negative skill gradient, especially among men. Second, distinguishing separations based

on the self-reported reason for being out of work, we find evidence consistent with the interpretation that recurrent separations stem from seasonal contractions in labor demand rather than from voluntary transitions. Third, we find that the large majority of excess recurrent separations represent repeat exits from the same employer, revealing that many seasonal work interruptions arise through iterative cycles of temporary layoffs followed by recall hiring. These facts provide important context for our subsequent analysis of household adaptation to seasonal work interruptions.

5.1 The skill gradient in excess recurrence

The left panel of [Figure 5](#) presents excess recurrence estimated separately by sex and by four education categories, using our primary SIPP sample. Among men, the likelihood of an annual recurrence is monotonically decreasing in educational attainment. Among women, the pattern is more mixed, with the highest rate observed among high school graduates and an uptick visible among college graduates as well. The high level of excess recurrence among college graduates is largely driven by frequent recurrent separations among college-educated teachers: as the middle panel shows, excluding workers whose initial separation was from educational services yields a skill gradient that is monotonically negative for men and nearly so for women.

Because job separations are more common among less-educated individuals, scaling our estimates by group-specific separation rates reinforces the impression of a negative skill gradient in exposure to labor market seasonality: the least-educated workers appear to be most susceptible to seasonal work interruptions, especially among men. The unequal prevalence of seasonal separations hints that seasonal volatility in earnings streams may be borne disproportionately by households with comparatively few financial resources, and it further motivates our later examination of household adaptation to such separations.

5.2 Reasons for unemployment

Seasonal work interruptions may have many underlying causes. In some cases, workers may be involuntarily separated from their employers at the end of the harvest cycle, construction schedule, tourist season, or school year. In other cases, workers may leave the labor force voluntarily at the same time each year, perhaps to pursue leisure opportunities or to attend to family responsibilities. Both conceptually and empirically, it is difficult to categorize particular exits from employment

Figure 5: Excess recurrence of job separations at 12 months: by sex $\times$ education

Notes: The left panel reports the excess recurrence of job separations at a 12-month horizon among prime-age SIPP respondents, obtained by estimating our preferred specification separately by sex and education level. The middle panel reports excess recurrence in a restricted sample that excludes workers whose initial separation was from educational services. The right panel again excludes the education sector, but scales each estimate by $12 \times$ the monthly unconditional separation rate among workers in the corresponding cell. Spikes show 95% confidence intervals, clustered by individual.

as either worker-initiated or employer-initiated because multiple factors may contribute to the dissolution of a given match. For example, a construction job may dissolve in the winter either because adverse weather interferes with on-site production or because construction workers find it costly to commute to work sites when road conditions are poor.

While it is impossible to draw a hard-and-fast line between voluntary quits and involuntary terminations, it is nonetheless informative to partition job separations on the basis of workers' self-reported reasons for being out of work. To do so, we leverage the fact that the CPS asks jobless workers who are currently looking for work why they are unemployed.¹⁸

Figure 6 reports estimates of excess recurrence for workers whose initial transition 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 job separation of *any kind*, including separations of types other than her 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

¹⁸We use our CPS sample to investigate this question since the SIPP is not well suited for this analysis. The CPS asks all unemployed workers to provide a reason for their unemployment, and these reasons are consistently coded across years. The SIPP asks some separating workers to provide a reason for their separation, but this question is not asked in every panel and, even when it is asked, many separators do not have a valid response reported.

Figure 6: 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-nonemployment transitions as recurrent separations, without regard for the type or reason for the second separation. Spikes show 95% confidence intervals, clustered by individual.

“job losers / on layoff” (estimate = 3.9 p.p., s.e. = 0.44), with small and statistically insignificant estimates for the remaining categories.

We view these patterns as consistent with an important role for labor demand factors, as evidenced by the fact that many seasonal separations stem from self-reported layoff. Even if they were undertaken willingly, however, seasonal jobless spells would still be likely to generate corresponding fluctuations in earned income, confronting households with the challenge of how to manage their expenditures in the face of volatile streams of income.

5.3 Repeat separations from the same employer

A venerable literature has documented that many job separations are in fact temporary layoffs, from which workers are ultimately recalled to their former employers. 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](#)),

Figure 7: 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 from the original employer, from a different employer in the original industry, or 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.

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*.¹⁹ Figure 7 decomposes the trajectory of recurrent separations for the 18 months following an initial separation into these three categories. Strikingly, we find that 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 periodic separations identified by our method and the layoff-recall phenomenon studied by prior literature.²⁰

¹⁹To 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.

²⁰Del Bono and Weber (2008) find that 64 percent 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 percent of new employment contracts separated by a temporary layoff are spaced 12 months apart, compared with 13 percent of contracts separated by a permanent layoff.

6 Earnings and Income Dynamics

Having established our method for identifying seasonal work interruptions, we now examine their consequences for earnings and household income. After laying out our econometric approach, we start by measuring the change in income associated with recurrent separations that are actually realized. We caution, however, that this exercise may incorrectly exclude those seasonal workers who promptly find alternative employment, while incorrectly including some nonseasonal workers who experience repeat separations purely by chance. To avoid such misclassifications, we construct a measure of individual exposure to seasonal work interruptions based on separators' excessive *propensity* to separate again 12 months later, as predicted using baseline characteristics. 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. Lost earnings from seasonal separations largely pass through to overall household income.

6.1 Econometric framework

Our goal is to estimate the change in income stemming from the ending of a seasonal job. For now, we focus on three primary outcomes of interest. First, we estimate the change in personal earnings, which represents the loss of income directly from the seasonal job net of any secondary jobs. Second, we estimate the change in household income, which captures the response of income along all margins at the time of seasonal job loss. 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 job loss. We defer to [Section 7](#) a look at the margins that account for this adaptation.

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

$$\mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 1], \quad \mathbb{E} [\Delta \text{Income}_{i,t} \mid \delta_{i,t} = 1], \quad \frac{\mathbb{E} [\Delta \text{Income}_{i,t} \mid \delta_{i,t} = 1]}{\mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 1]}, \quad (3)$$

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 a proxy for seasonal job loss. We consider two different methods for constructing this proxy, an indicator variable that we denote by $\hat{\delta}_{i,t}$. The first method, which is based on *realized* recurrent separations, is simple and easy to interpret, but it could potentially exclude important types of seasonal separators. The second is based on *probable* recurrent separations and is able to include seasonal separators who are excluded by the first method. We describe both of these methods for constructing $\hat{\delta}_{i,t}$ in detail in subsequent sections.

Once we have constructed a proxy indicator $\hat{\delta}_{i,t}$ for seasonal job loss 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 from [Equation 3](#), i.e.,

$$\mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \hat{\delta}_{i,t} = 1] \approx \mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 1] \quad (4)$$

We can compute approximations for the other objects of interest from [Equation 3](#) in the same fashion, by conditioning the expectations on the proxy indicator $\hat{\delta}_{i,t}$ in place of the unobserved indicator of seasonal job loss $\delta_{i,t}$.

To differing degrees, both of our proxy indicators may include some nonseasonal 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 job loss (i.e., $\hat{\delta}_{i,t} = 1$), a fraction p of these observations represent actual seasonal job loss ($\delta_{i,t} = 1$), while a fraction $(1 - p)$ represent nonseasonal workers ($\delta_{i,t} = 0$). Our estimate of the change in earnings will be a mixture of the changes for the two groups,

$$\mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \hat{\delta}_{i,t} = 1] = p \mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 1] + (1 - p) \mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 0] \quad (5)$$

with the same identity also holding with $\text{Income}_{i,t}$ in place of $\text{Earnings}_{i,t}$. If nonseasonal workers are

in steady-state, with no changes in earnings or income on average ($\mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 0] = 0$, $\mathbb{E} [\Delta \text{Income}_{i,t} \mid \delta_{i,t} = 0] = 0$), then the change in each component is attenuated by the fraction of nonseasonal workers, $1 - p$. Importantly, even though these estimates of the changes are attenuated, our estimate of the pass-through rate will recover the true pass-through rate without attenuation:

$$\frac{\mathbb{E} [\Delta \text{Income}_{i,t} \mid \hat{\delta}_{i,t} = 1]}{\mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \hat{\delta}_{i,t} = 1]} = \frac{p \mathbb{E} [\Delta \text{Income}_{i,t} \mid \delta_{i,t} = 1]}{p \mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 1]} = \frac{\mathbb{E} [\Delta \text{Income}_{i,t} \mid \delta_{i,t} = 1]}{\mathbb{E} [\Delta \text{Earnings}_{i,t} \mid \delta_{i,t} = 1]} \quad (6)$$

The pass-through rate is still informative even if nonseasonal workers are not in steady-state. In [Appendix C](#) we derive conditions under which the pass-through rate that we measure empirically forms a lower bound for the true pass-through rate among seasonal workers. These conditions are likely to hold *a priori*, and we also present supporting evidence in the appendix. We thus interpret our estimates of pass-through as representing a lower bound for the true pass-through rate.

6.2 Income dynamics of recurring separators

We begin by proxying for seasonal job loss using *realized* annually recurrent separations. For a sample of N individuals who each experience two separations spaced 12 months apart, 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

$$\beta_y^{(\tau)} \equiv \frac{1}{N} \sum_{i=1}^N \frac{y_{i,t_0+\tau} - y_{i,t_0+\tau-1}}{\text{Earnings}_{i,t_0-1}} \quad (7)$$

Thus $\beta_y^{(\tau)}$ is the average month-over-month change in y among recurrent separators, with changes denominated by the separator's earnings one month prior to the initial separation.²¹ 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 preseparation 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 percent and 300 percent of the focal separator's

²¹Paralleling our approach in [Section 4.1](#), we stack a copy of individual i 's data for each separation she experiences, so that we avoid arbitrarily treating some separations as base separations and others as recurrences.

preseparation earnings.²² We adjust changes in earnings and (later) other components of income on a proportionate basis to ensure that they add up 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:

$$\tilde{\beta}_y^{(n)} \equiv \beta_y^{(0)} + \beta_y^{(1)} + \dots + \beta_y^{(n)} \quad (8)$$

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 nonrandom attrition. Taking long differences yields similar results.

Note that the coefficients $\tilde{\beta}_y^{(n)}$ do not represent the *percent change* in the outcome y : instead, they are the change in y *as a percentage of preseparation earnings*, which enables more straightforward comparisons. For example, consider a household composed of a seasonal worker who earns \$1,000 per month and a nonseasonal 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 loses her job and this income is not replaced by any other source, so that her earnings fall by \$1,000. This is both a 100 percent decrease and a decrease as a percent of preseparation earnings of 100 percentage points (p.p.). However, although household income has only fallen by 25 percent, the \$1,000 decrease in household income *as a percentage of the seasonal worker's preseparation 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 comparison of the percent changes in these variables.

Applying this procedure, [Figure 8](#) shows the estimated changes in total personal earnings for individuals who separate from employment on two occasions exactly 12 months apart (left panel), along with changes in household income (right panel). Because both variables are normalized by preseparation 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 changes of identical magnitude.

²²Our \$450 threshold for (deflated) preseparation 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.

Figure 8: Earnings and household income of recurrent separators

Notes: The left panel plots the evolution of personal earnings for prime-age SIPP respondents who experience two separations spaced exactly 12 months apart. The 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. The monthly change in household income is winsorized to be between -300% and 300% , and the monthly change in earnings is adjusted proportionally for winsorized observations. Individuals with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Spikes show 95% confidence intervals, clustered by individual.

The volatile earnings of recurring separators induce large swings in household income. Earnings fall after the initial separation, almost entirely recover over the course of the subsequent year, and then plummet again at the time of the second separation. Personal earnings as a percent of pre-separation earnings fall by 48.0 p.p. between event-times $\tau = 10$ and $\tau = 13$ months, and the large majority of this lost income passes through to household income, which (in comparable units) falls by 37.5 p.p. over the same time period. The ratio of these two declines implies that only a modest portion of lost earnings—\$0.22 for each \$1 lost—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 nonlabor sources.

However, the earnings and income dynamics seen among recurrent separators may not be representative of seasonal workers as a whole. By construction, this sample excludes any individual who held a seasonal job, anticipated an impending separation at $\tau = 12$ months, and accordingly secured a new job to begin immediately after the seasonal job ended. By conditioning on an observed interval of joblessness occurring one year after an earlier separation, we therefore shut

down a prime margin along which workers might adapt to seasonal work interruptions. A separate concern is that this sample will inevitably include some individuals who, though not employed in seasonal work, experienced annually recurrent separations purely by chance. To mitigate these misclassification errors, we next isolate a set of workers whose initial separations bear the hallmarks of seasonal job loss, whether or not they separated again the following year.

6.3 Pinpointing seasonal separators

To address the potential problems with the sample discussed above, we construct an alternative proxy indicator for seasonal job loss from a sample of separators who appear *likely* to be engaged in seasonal work, and hence *likely* to separate again at the onset of their idiosyncratically timed “off-seasons.” To do so, we estimate heterogeneity in excess recurrence across job types and use these estimates to predict each separator’s likelihood of experiencing another separation 12 months after her first one. 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 job separators by our measure of their propensity to separate again, we find that those in the top decile of our measure have an average excess recurrence rate about four times higher than those in the bottom decile.

We start by estimating the variation in recurrent separations across different types of jobs. For a worker in job 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 $\rho_\tau(j) \equiv \Pr(\text{Sep}_{t_0+\tau} \mid \text{Sep}_{t_0}, X_j)$. We treat this probability as an unknown (possibly nonlinear) 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 $\tilde{\rho}(j) \equiv f(X_j, \tau = 12) - \frac{1}{2} (f(X_j, \tau = 11) + f(X_j, \tau = 13))$.

One approach for estimating 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 nonparametric, 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 nonparametrically using gradient-boosted trees, a widely used machine-learning algorithm for estimating high-dimensional functions with unknown functional form.²³ 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 B for details on the temperature data). Lastly, we include the number of months since the initial separation τ as well as the calendar month (January, February, ...) in which the original separation occurred.²⁴ We emphasize that gradient-boosted trees naturally allow for flexible interactions among the input variables, without the econometrician needing to prespecify which interactions ought to be included in the model.

We estimate the function f using our CPS sample.²⁵ 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.²⁶ Minimizing cross-validated mean squared error

²³This algorithm was developed by Breiman (1997) and Friedman (2001, 2002). Friedman et al. (2001) provide a helpful overview for readers unfamiliar with the method.

²⁴Industry and occupation are treated as categorical variables, while January 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. The package we use to implement gradient-boosted trees, LightGBM, automatically converts categorical variables to their one-hot encoded representation during estimation (Ke et al., 2017).

²⁵By 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.

²⁶For a single 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 possible combinations of the number of trees in the set {50, 100, 200, 400, 600, 800, 1000}

Figure 9: 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, as measured in the National Oceanic and Atmospheric Administration’s nCLIMDIV database. Decile cutoffs incorporate sampling weights.

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 initially in job j we construct the predicted excess recurrence rate $\hat{\rho}_i$ as:

$$\hat{\rho}_i = \hat{f}(X_j, \tau = 12) - \frac{1}{2} \left(\hat{f}(X_j, \tau = 11) + \hat{f}(X_j, \tau = 13) \right) \quad (9)$$

This forms an estimate of the extent to which we would predict an individual who is separating from a job today is excessively likely to separate from a job 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.

To illustrate the kind of subtle interactions that our algorithm is able to detect, [Figure 9](#) shows the share of construction workers whom our algorithm assigns to the top decile of predicted excess

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 in each leaf.

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.²⁷

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 D.2](#) shows that this is indeed the case: the estimated average excess recurrence rate is 6.1 percent for those in the top decile, about four times larger than the average rate of 1.5 percent for those in the bottom decile.

6.4 Income dynamics of predicted separators

Armed with our ranking of workers by predicted excess recurrence, we return to estimating the evolution of income around the onset of seasonal separations. [Figure 10](#) displays the evolution of personal earnings and household income for individuals who initially separate from a job and are in either the top or bottom decile of predicted excess recurrence. As previously, both outcomes are reported as a percentage of preseparation earnings in order to aid comparison.

The estimates shown in [Figure 10](#) 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, earnings decline by 18.6 p.p. between 10 and 13 months after the initial separation; household income as a percent of preseparation earnings declines 17.3 p.p. over the same period. This amounts to a decline in household income of about \$0.93 for every \$1 of lost earnings from a seasonal job. This pattern for the most seasonal separators stands in contrast to that of the least seasonal (bottom-decile) separators, who experience a 2.5 p.p. increase in earnings and a 3.3 p.p. increase in household income over the same time period.

²⁷ [Appendix Table D.2](#) reports several characteristics of separators classified as seasonal by our algorithm. Throughout the paper, we use sampling weights when calculating decile cutoffs, to ensure that the top and bottom deciles each represent one-tenth of the population of job separations.

Figure 10: 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 (“most seasonal”) or bottom (“least seasonal”) decile 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. The monthly change in household income is winsorized to be between -300% and 300% and the monthly change in earnings is adjusted proportionally for winsorized observations. Individuals with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Spikes show 95% confidence intervals, clustered by individual.

The difference in evolution of income for the most and least seasonal separators highlights the importance of using predicted excess recurrence to identify seasonal separations. The analysis in [Section 6.2](#) limited the sample to those individuals with realized separations in both months 0 and 12, thereby excluding seasonal workers who were able to avoid a jobless spell at month 12. The approach presented here avoids this misclassification, as the “most seasonal” category in [Figure 10](#) includes this group. Additionally, as a result of including individuals who are predicted to experience a recurrent seasonal separation but ultimately do not, the decline in earnings and income 12 months after the initial separation is smaller for the most seasonal separators in [Figure 10](#) than for the separators in [Figure 8](#). It is, though, larger than the 6.1 percent excess recurrence rate for the most seasonal separators shown in [Appendix Figure D.2](#), indicating that some workers who do not experience a transition from work into nonwork nonetheless experience a decline in earnings.

These comparisons relate to earnings and income *received in month* $\tau = 13$ relative to month $\tau = 10$. To get an estimate of households’ *cumulative* 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 $\Delta Y_i \equiv \sum_{\tau=11}^{13} \frac{y_{i,t_0+\tau} - y_{i,t_0+10}}{\text{Earnings}_{i,t_0-1}}$. In [Appendix Figure D.3](#), we present the empirical cdf of ΔY_i for households in the bottom and top deciles of predicted excess recurrence. Among the most seasonal households, 28.1 percent face cumulative off-season income losses exceeding one month of the separator’s preseparation earnings. Roughly one in six seasonal households face cumulative 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.

7 Margins of Adaptation to Seasonal Work Interruptions

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 the individuals most and least prone to exhibiting a recurrent job 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.

7.1 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 3](#) 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,

Table 3: Changes in separator earnings by exposure to seasonal work interruptions

All SIPP panels			1990–1993 panels only		
	Most seasonal	Least seasonal		Most seasonal	Least seasonal
Total earnings	-18.62 (1.04)	2.52 (0.99)	Total earnings	-24.92 (2.09)	3.62 (1.83)
<i>Original industry</i>	-19.80 (0.90)	-1.01 (0.81)	<i>Original employer</i>	-22.69 (1.68)	-2.75 (1.23)
			<i>Other employers in original industry</i>	-2.82 (1.12)	2.72 (1.48)
<i>Other industries</i>	1.18 (0.61)	3.53 (0.72)	<i>Other industries</i>	0.58 (1.03)	3.65 (1.06)

Notes: Changes in components of separators’ earnings from 10 to 13 months after an initial separation for SIPP separators in the top (“most seasonal”) and bottom (“least seasonal”) deciles of the predicted excess recurrence distribution. The right panel uses the 1990–1993 SIPP sample with corrected job identifiers described in the notes for Figure 7. All variables are expressed as percentages of the individual’s earnings one month prior to the initial separation. The components are adjusted on a monthly basis to ensure that the total household income change in any month is between -300% and 300% , consistent with the winsorized estimates presented previously. We exclude individuals with pre-separation earnings less than \$450/month in January 2017 dollars. Standard errors are shown in parentheses.

with these groups defined as in the previous section.

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.²⁸ Earnings growth in other industries, though positive, is quantitatively 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” job opportunities pending the next seasonal upswing.

7.2 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

²⁸As 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.

Figure 11: Changes in partner earnings, decomposed by industry

Notes: Changes in partner's earnings by industry for SIPP separators in the top ("most seasonal") and bottom ("least seasonal") deciles 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. The earnings components are adjusted on a monthly basis to ensure that the total household income change in any month is between -300% and 300% , consistent with the winsorized estimates presented previously. Partners of separators with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Spikes show 95% confidence intervals, clustered by individual.

supply (Lundberg, 1985; Spletzer, 1997; Stephens, 2002; Juhn and Potter, 2007). 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 11 shows that partner earnings for the 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 a decrease of 0.1 p.p. and an increase of 1.2 p.p., respectively, for the least seasonal separators (see Appendix Table D.3). 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. 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.²⁹ For seasonal workers, having a partner who works in the same industry creates correlated income volatility over the year, diminishing the ability of this margin to smooth out one’s own income fluctuations.

The difference between this result and the previous literature appears to be due to our focus on seasonal separators, and not to differences in our approach or methodology. As seen in [Figure 11](#), 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 D.3](#) reports changes in the earnings of other, nonpartner 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 separation, indicating that this margin does not much mitigate or exacerbate lost earnings from seasonal work interruptions.

7.3 Government transfers

A final margin through which households may replace lost earnings is the collection of government transfers. We focus on three types of transfers in particular: 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 U.S. 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 preseparation earnings rise by 4.7 p.p. among the most seasonal separators between 10 and 13

²⁹Furthermore, [Hyatt \(2019\)](#) shows that many couples who work in the same industry in fact work at the same establishment, potentially increasing the degree to which earnings are correlated for couples working in seasonal jobs.

Figure 12: Changes in receipt of unemployment benefits, adjusted for underreporting

Notes: Changes in household UI income for SIPP separators in the top (“most seasonal”) and bottom (“least seasonal”) deciles of the predicted excess recurrence distribution. We scale up UI income at the individual level using the year-by-year estimates reported in [Meyer et al. \(2015\)](#) 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. UI income is adjusted on a monthly basis to ensure that the total household income change in any month is between -300% and 300% , consistent with the winsorized estimates presented previously. Households containing separators with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Spikes show 95% confidence intervals, clustered by individual.

months after an initial separation, compared to no change for the least seasonal separators.³⁰ This offsets approximately one-quarter of lost earnings.

Although the SIPP is “specifically designed to determine eligibility and receipt of government transfers” ([Meyer et al., 2015](#)), households may nonetheless underreport the receipt of government transfers. [Meyer et al.](#) 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 percent 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 et al. \(2015\)](#) and recompute all of our estimates. Adjusted for underreporting, [Figure 12](#) shows that UI receipts rise by 6.9 p.p. among the most seasonal separators, offsetting about one-third of lost earnings.

In contrast, [Appendix Table D.4](#) 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

³⁰Seasonal separators also appear to receive more UI benefits than nonseasonal separators during their initial separation, possibly indicating that UI is an implicit part of the labor contract in seasonal sectors ([Feldstein, 1975](#)).

0.6 p.p. increase in receipt of SNAP benefits among the most seasonal households, relative to a base receipt rate of 10.1 percent for this group. All in all, however, the patterns we observe are consistent with UI providing the main source of social insurance for seasonal separations.

7.4 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: 1) own earnings, 2) partner earnings, 3) other co-resident earnings, 4) household UI income, 5) other household transfer income, and 6) all other sources. By normalizing each of these components relative to preseparation 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 4](#) shows the full decomposition of the change in household income between 10 and 13 months after an initial separation for both the most and least 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 spousal earnings. Other categories have comparatively little influence on the bottom line. In contrast, the least seasonal separators see increases in almost all components of household income (see [Appendix Table D.5](#)).

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 is a lower bound. Though we have limited our sample to separations characteristic of seasonal work interruptions, we have presumably misclassified some nonseasonal separators as seasonal. As detailed in [Appendix C](#), our empirical results imply that such misclassification biases our pass-through estimate downward, so that true pass-through is higher.

In total, 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 unsuitable for seasonal separators. During a seasonal dip in the labor market, it will be relatively difficult to find a new job or have a spouse go back to work. At the same time, many seasonal separators will have earned enough to qualify for unemployment insurance, and so it should not be surprising that this is the main margin along which they do smooth income.

Table 4: Decomposing changes in household income among probable seasonal separators

	Overall	Sex of focal separator:		Type of separation:		Pre-separation sector:	
		Male	Female	Unemp.	Non-part.	Education	Other
<i>A. Unadjusted estimates</i>							
Own earnings	-18.62 (1.04)	-17.81 (1.52)	-19.42 (1.42)	-22.56 (1.29)	-11.22 (1.71)	-22.91 (2.02)	-16.96 (1.21)
Partner earnings	-2.60 (0.97)	-2.07 (0.75)	-3.15 (1.82)	-1.94 (0.93)	-3.86 (2.14)	-1.13 (2.37)	-3.18 (1.00)
Other co-resident earnings	-1.13 (0.81)	-3.43 (1.10)	1.23 (1.18)	-1.95 (0.94)	0.42 (1.48)	4.00 (1.55)	-3.04 (0.94)
Household UI benefits	4.65 (0.40)	6.34 (0.60)	2.91 (0.54)	5.97 (0.52)	2.17 (0.59)	0.19 (0.40)	6.32 (0.53)
Other transfer receipts	0.21 (0.13)	0.20 (0.20)	0.21 (0.17)	0.14 (0.18)	0.33 (0.19)	0.10 (0.11)	0.25 (0.18)
All other sources	0.15 (0.50)	-0.06 (0.65)	0.37 (0.76)	0.60 (0.56)	-0.70 (0.98)	0.54 (1.17)	0.01 (0.53)
Household income	-17.34 (1.55)	-16.83 (1.88)	-17.83 (2.48)	-19.75 (1.76)	-12.87 (2.97)	-19.22 (3.45)	-16.61 (1.70)
Pass-through rate	93.1 (6.8)	94.5 (8.0)	91.8 (11.0)	87.5 (6.1)	114.7 (23.1)	83.9 (12.5)	97.9 (8.3)
<i>B. Adj. for UI underreporting</i>							
Household UI benefits	6.93 (0.60)	9.29 (0.88)	4.49 (0.80)	8.80 (0.76)	3.40 (0.88)	0.34 (0.57)	9.39 (0.78)
Household income	-15.06 (1.55)	-13.87 (1.87)	-16.26 (2.48)	-16.92 (1.75)	-11.64 (2.97)	-19.06 (3.46)	-13.54 (1.69)
Pass-through rate	80.9 (7.0)	77.9 (8.4)	83.7 (11.1)	75.0 (6.4)	103.7 (23.1)	83.2 (12.6)	79.8 (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. The change in household income is winsorized on a monthly basis to be between -300% and 300% of pre-separation earnings, and each of the components of household income is adjusted proportionately on a monthly basis to ensure that they add up to the winsorized total. Individuals with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Standard errors shown in parentheses.

7.5 Differences in adaptation across separators

Next we turn to examining whether particular types of seasonal separations drive our main results. One possible concern with the results shown previously is that they may be driven by a particular subgroup and hence not representative of all seasonal separations. We repeat the additive decomposition of household income outlined in the previous section separately for three different cuts of the sample: splitting 1) by gender, 2) by labor force status after the initial separation, and

3) by initial sector (educational services or all others). [Table 4](#) shows the decomposition for the most seasonal separators in each subgroup.³¹ As above, all changes are normalized by the focal separator’s preseparation earnings.

We find small differences in our main results across each of the subgroups we examine. Male and female seasonal separators exhibit similar declines in own earnings between 10 and 13 months after their original separation (17.8 p.p. and 19.4 p.p., respectively), as well as similar declines in partner earnings (2.1 p.p. and 3.2 p.p.). Separations into unemployment exhibit larger earnings declines than separations into nonparticipation (22.6 p.p. vs. 11.2 p.p.) but smaller declines in partner earnings (1.9 p.p. vs. 3.9 p.p.). Both educators and noneducators see large declines in own earnings (22.9 p.p. and 17.0 p.p., respectively), but educators see a smaller decline in partner earnings (1.1 p.p. vs. 3.2 p.p.). Household UI income grows more for men than women, unemployed separators than nonparticipants, and noneducators than educators.

Across all subgroups, though, 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. The adjusted pass-through rate is slightly larger for women than men (84 percent vs. 78 percent), larger for nonparticipants than unemployed separators (104 percent vs. 75 percent), and larger for educators than noneducators (83 percent vs. 80 percent). While this heterogeneity may be indicative of differing capacities to make up for lost earnings across these different groups, the consistently high pass-through rates indicate that an inability (or unwillingness) to make up for lost earnings is widespread among seasonal workers.

8 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.

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 recov-

³¹[Appendix Table D.5](#) shows the same set of estimates for the least seasonal separators.

ery 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.

Seasonal separations are related to other types of job separations studied in the economics literature, but they are distinct in some important respects. The pattern of regular periods of work, punctuated by temporary layoffs and recalls, is reminiscent of the workers studied by [Feldstein \(1975\)](#). By contrast, the fact that seasonal workers’ earnings return to their previous levels when they return to work is unlike the scarring effect created by mass layoffs ([Jacobson et al., 1993](#); [Davis and von Wachter, 2011](#)). Additionally, whereas other types of job separations tend to see an *added* worker effect, seasonal separators exhibit a *subtracted* worker effect ([Lundberg, 1985](#); [Spletzer, 1997](#); [Stephens, 2002](#); [Juhn and Potter, 2007](#)).

The patterns documented in this paper are particularly important because other forms of income volatility intersect with seasonal work. Many previous studies that estimate total income volatility from all sources have relied on annual measures of earnings to measure the magnitude of income fluctuations in the U.S. labor market ([Blundell et al., 2008](#); [Sabelhaus and Song, 2009](#); [Kopczuk et al., 2010](#); [DeBacker et al., 2013](#)). But abstracting from within-year sources of income volatility, such as the variable earnings garnered in seasonal work, understates the degree of turbulence in many households’ financial lives. At the other extreme, changes in high-frequency income dynamics are layered on top of preexisting seasonal patterns. For instance, recent research has examined the rise of alternative work arrangements, which have created new forms of short-term employment for many workers ([Farrell and Greig, 2015](#)). The growth of the gig economy may create new categories of seasonal work, but depending on how such work is distributed throughout the year, it could at the same time dampen seasonal fluctuations in household income by giving already-seasonal workers increased opportunities to earn income in the off-season.

Finally, 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, recently proposed new 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 et al., 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.
- Beaulieu, J. Joseph, Jeffrey K. MacKie-Mason, and Jeffrey A. Miron. 1992. "Why Do Countries and Industries with Large Seasonal Cycles Also Have Large Business Cycles?" *The Quarterly Journal of Economics*, 107(2): 621–656.
- Beaulieu, J. Joseph, and Jeffrey A. Miron. 1992. "A Cross Country Comparison of Seasonal Cycles and Business Cycles." *The Economic Journal*, 102(413): 772–788.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.
- Breiman, Leo. 1997. "Arcing the Edge." Technical Report 486, Statistics Department, University of California at Berkeley.
- Browning, Martin, and Thomas F Crossley. 2001. "Unemployment Insurance Benefit Levels and Consumption Changes." *Journal of Public Economics*, 80(1): 1–23.
- Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *The Quarterly Journal of Economics*, 112(1): 1–55.
- Couch, Kenneth A., and Dana W. Placzek. 2010. "Earnings Losses of Displaced Workers Revisited." *American Economic Review*, 100(1): 572–589.
- Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Cost of Job Loss." *Brookings Papers on Economic Activity*, Fall, 1–72.
- de Raaf, Shawn, Costa Kapsalis, and Carole Vincent. 2003. "Seasonal Work and Employment Insurance Use." *Perspectives on Labour and Income*, 4(9): 5–11.

- 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.
- 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 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*. 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. by Arnold Zellner: Cambridge, MA: National Bureau of Economic Research, 30–32.
- 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. 2019. "Job Displacement Insurance and (the Lack of) Consumption-Smoothing." NBER Working Paper 25749.
- Geremew, Menelik, and François Gourio. 2018. "Seasonal and Business Cycles of U.S. Employment." *Economic Perspectives*, 42(3): 1–28.
- 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.
- Juhn, Chinhui, and Simon Potter. 2007. "Is There Still an Added Worker Effect?" Federal Reserve Bank of New York Staff Report No. 310.
- Kalton, Graham, Marianne Winglee, Louis Rizzo, Thomas Jabine, and Daniel Levine. 1998. "SIPP Quality Profile 1998." SIPP Working Paper 230, US Census Bureau.
- 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.
- Katz, Lawrence F., and Bruce D. Meyer. 1990. "Unemployment Insurance, Recall Expectations, and Unemployment Outcomes." *The Quarterly Journal of Economics*, 105(4): 973–1002.
- Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. "LightGBM: A Highly Efficient Gradient Boosting Decision Tree." In *Advances in Neural Information Processing Systems*. 3146–3154.
- Kiefer, Nicholas M. 1988. "Economic Duration Data and Hazard Functions." *Journal of Economic Literature*, 26(2): 646–679.
- Kopczuk, Wojciech, Emmanuel Saez, and Jae Song. 2010. "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data Since 1937." *The Quarterly Journal of Economics*, 125(1): 91–128.
- Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury. 2020. "Sources of Displaced Workers' Long-Term Earnings Losses." *American Economic Review*, forthcoming.
- 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." In *Advances in Neural Information Processing Systems*. 512–518.
- 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).
- Miron, Jeffrey A., and J. Joseph Beaulieu. 1996. "What Have Macroeconomists Learned about Business Cycles from the Study of Seasonal Cycles?" *Review of Economics and Statistics*, 78(1): 54–66.
- Miron, Jeffrey A., and Stephen P. Zeldes. 1988. "Seasonality, Cost Shocks, and the Production Smoothing Model of Inventories." *Econometrica*, 56(4): 877–908.

- 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 Semi-parametric 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.
- Sabelhaus, John, and Jae Song. 2009. "Earnings Volatility Across Groups and Time." *National Tax Journal*, 62(2): 347–364.
- 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.
- Spletzer, James R. 1997. "Reexamining the Added Worker Effect." *Economic Inquiry*, 35(2): 417–427.
- 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 the 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.

Appendix A 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 nonseasonal 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 her “off-season,” so that she experiences exactly one hire and one separation each year. By contrast, each nonseasonal worker has stochastic employment governed by constant one-month job-finding and separation hazard rates f_i and s_i , so that her 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 $\Pr_i(\text{Sep}_t)$ be the probability that worker i separates from a job in month t , i.e., that she 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 $\Pr_i(\text{Sep}_{t+\tau} \mid \text{Sep}_t)$. Nonseasonal workers are approximately as likely to separate from a new job either 11, 12, or 13 months after a previous separation.³² 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

$$\text{excess recurrence} \equiv \Pr(\text{Sep}_{t+12} \mid \text{Sep}_t) - \frac{1}{2} (\Pr(\text{Sep}_{t+11} \mid \text{Sep}_t) + \Pr(\text{Sep}_{t+13} \mid \text{Sep}_t)) \quad (10)$$

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

³²For a nonseasonal 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 $\tau \in \{11, 12, 13\}$.

can be expressed in terms of α :

$$\text{excess recurrence} \approx \frac{\frac{\alpha}{12}}{(1 - \alpha) \int \frac{s_i f_i}{s_i + f_i} di + \frac{\alpha}{12}} \quad (11)$$

where the denominator is equal to the average monthly separation rate in the economy, $\text{Pr}(\text{Sep})$.³³

Intuitively, excess recurrence represents the fraction of separations in the average month which come from seasonal workers. Multiplying this quantity by the average separation rate and then multiplying by 12 gives the total proportion of seasonal workers in this economy:

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

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.

This stylized model makes a number of simplifying assumptions to derive the identity in Equation 12, 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 nonseasonal 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.

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 B Additional Details on Sample Construction

Our analysis draws on two workhorse U.S. household survey datasets, the Survey of Income and Program Participation (SIPP) and the Current Population Survey (CPS), as well as climatological data from the National Oceanographic and Atmospheric Administration (NOAA).

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

³³The approximation error in this expression arises because, for generic values of s_i and f_i , the separation rate ρ_τ among nonseasonal 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.

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).³⁴ 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 nonresponse that the Census Bureau populates through imputation. Second, we restrict to ages 25–54, discarding all subsequent observations once an individual turns 55. 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 U.S. population at each point in time.

With this sample in hand, we harmonize the demographic, employment, and income variables needed for our analysis. As described in the main text, we use the underlying week-level detail on labor force status to construct an indicator variable for experiencing at least one job separation in a given month, and to distinguish separations into unemployment from those into nonparticipation. Following standard practice, 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).³⁵ For comparability between the SIPP and the CPS, we aggregate the SIPP’s detailed industry codes into 15 one-digit sectors, augmented with a missing category. We similarly aggregate the SIPP occupational codes into 14 one-digit occupations plus a missing category. Finally, we deflate all monetary amounts to January 2017 dollars using the Federal Reserve Board’s chain-type personal consumption expenditures price index.

Except where indicated, we employ an unbalanced panel throughout our analysis, with individuals retained in-sample as long as possible to maximize statistical power.

CPS data

We use monthly core CPS files provided by the Integrated Public Use Microdata Series, or IPUMS (Flood et al., 2018). Paralleling our SIPP sample period, we stack CPS files 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. As with the SIPP, we restrict our sample to the civilian population ages 25–54.

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 either 1) she was at work or 2) she

³⁴The 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>.

³⁵A small number of observations are missing education codes. Since educational attainment is likely to be time-invariant in our prime-age population, we fill missing values where possible by forward and backward extrapolation, and we recode the few remaining missings (less than 0.1 percent of the sample) to the lowest educational category.

held a job but was not at work during the reference week. We code a worker as experiencing a job separation if she was employed in the previous month but not in the current month.

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 seven-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. Lacking data on discretionary shifts in timing, we apply the automatic portion of 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.

Beginning with the IPUMS variables EDUC, IND1990, and OCC1990 (which have been harmonized over time), we recode education, industry, and occupation to the same categories described above for our SIPP sample. We weight respondents by the final person-level weights provided by IPUMS, and we again deflate monetary amounts to January 2017 dollars.

As with the SIPP, we employ an unbalanced CPS panel throughout our analysis, with individuals retained in-sample for as many observations as possible.

NOAA data

To measure each state’s typical January temperature, we compute the average maximum daily temperature for the Januaries of 1984–2013 using state-level data from NOAA’s nCLIMDIV 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 SIPP sometimes combines certain less-populous states when reporting state of residence (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 C Bounding the Pass-through Rate

Using the notation of [Section 6.1](#) (suppressing subscripts throughout for simplicity), we can express the expected changes in earnings and household income among observations we classify as seasonal job losses—that is, those for which our proxy $\hat{\delta}$ equates to 1—as mixtures of these changes for seasonal and nonseasonal workers:

$$\begin{aligned}\mathbb{E} \left[\Delta \text{Earnings} \mid \hat{\delta} = 1 \right] &= p \mathbb{E} [\Delta \text{Earnings} \mid \delta = 1] + (1 - p) \mathbb{E} [\Delta \text{Earnings} \mid \delta = 0] \\ \mathbb{E} \left[\Delta \text{Income} \mid \hat{\delta} = 1 \right] &= p \mathbb{E} [\Delta \text{Income} \mid \delta = 1] + (1 - p) \mathbb{E} [\Delta \text{Income} \mid \delta = 0]\end{aligned}\tag{13}$$

where p is the fraction of observations with true seasonal job loss.

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

nonseasonal workers (θ^{NS}):

$$\begin{aligned}\theta^M &\equiv \frac{\mathbb{E}[\Delta \text{Income} \mid \hat{\delta} = 1]}{\mathbb{E}[\Delta \text{Earnings} \mid \hat{\delta} = 1]} \\ \theta^S &\equiv \frac{\mathbb{E}[\Delta \text{Income} \mid \delta = 1]}{\mathbb{E}[\Delta \text{Earnings} \mid \delta = 1]} \\ \theta^{NS} &\equiv \frac{\mathbb{E}[\Delta \text{Income} \mid \delta = 0]}{\mathbb{E}[\Delta \text{Earnings} \mid \delta = 0]}\end{aligned}\tag{14}$$

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:

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

which simplifies to

$$\theta^M = \theta^S + \underbrace{\frac{(1-p)(\theta^{NS} - \theta^S)\mathbb{E}[\Delta \text{Earnings} \mid \delta = 0]}{\mathbb{E}[\Delta \text{Earnings} \mid \hat{\delta} = 1]}}_{\text{bias term}}\tag{16}$$

The sign of the bias term depends on its components. Since $\hat{\delta}$ is a proxy for seasonal job loss, which lowers 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 10](#) drifts upwards steadily throughout the 18 months following an initial separation.

Therefore, the sign of the bias term depends inversely on the sign of $(\theta^{NS} - \theta^S)$. 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 nonseasonal 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 10](#), where workers predicted to be nonseasonal have pass-through greater than 100 percent while workers predicted to be seasonal have pass-through less than 100 percent.

Appendix D Additional Results

Figure D.1: Excess recurrence by month of separation, scaled by separation rates

Notes: Excess annual recurrence of job separations in our SIPP sample, obtained by estimating our preferred specification separately by the month in which the original separation occurred and then scaling excess recurrence by $12 \times$ the monthly separation rate in that month. Spikes show 95% confidence intervals, clustered by individual.

Figure D.2: 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, as described in [Section 4.1](#). Spikes show 95% confidence intervals, clustered by individual. Decile cutoffs incorporate sampling weights.

Figure D.3: The cdf of cumulative off-season shortfalls in household income

Notes: The sample consists of prime-age SIPP job separators who are in either the top (“most seasonal”) or bottom (“least seasonal”) decile of predicted excess recurrence. For each household i , we calculate the total off-season shortfall 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.

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

	Employed workers		All job separators		Recurrent separators	
	(N = 15,383,408)		(N = 300,584)		(N = 7,096)	
<hr/>						
Demographics <hr/>						
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)
<hr/>						
Household structure <hr/>						
Married	66.9	(47.1)	61.2	(48.7)	68.0	(46.6)
≥1 children in household	54.2	(49.8)	57.3	(49.5)	61.2	(48.7)
<hr/>						
Select industry indicators <hr/>			<i>(For separators: measured pre-separation)</i>			
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)

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 nonemployed 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).

Table D.2: Summary statistics: SIPP separators least/most likely to separate again

	Least seasonal		Most seasonal	
	(N = 12,053)		(N = 12,678)	
<hr/>				
Demographics				
<hr/>				
Female	29.8	(45.7)	47.8	(50.0)
Age	36.7	(8.3)	37.6	(8.4)
Non-white	35.4	(47.8)	26.8	(44.3)
Non-college	59.3	(49.1)	51.6	(50.0)
<hr/>				
Household structure				
<hr/>				
Married	58.9	(49.2)	64.2	(47.9)
≥1 children in household	56.3	(49.6)	57.9	(49.4)
<hr/>				
Pre-separation industries (select examples)				
<hr/>				
Agriculture, fishing, & forestry	1.2	(10.7)	12.0	(32.5)
Construction	38.1	(48.6)	14.7	(35.4)
Educational services	9.4	(29.1)	27.6	(44.7)
Healthcare	2.3	(15.0)	0.3	(5.5)
<hr/>				
Pre-separation monthly earnings (2017\$)				
<hr/>				
Personal earnings	2,774.7	(2786.9)	2,532.6	(2192.4)
Household income	5,330.6	(4297.3)	5,838.0	(4535.9)

Notes: All columns restrict to job separators ages 25–54 observed during 1984–2013. “Least seasonal”: restrict to separators in the bottom decile 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 D.3: 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 (0.97)	1.14 (0.76)	-4.68 (1.38)	0.66 (1.13)
<i>... in separator's industry</i>	-2.29 (0.48)	-0.07 (0.26)	-3.25 (0.66)	-0.12 (0.39)
<i>... in different industries</i>	-0.31 (0.90)	1.21 (0.75)	-1.43 (1.29)	0.78 (1.11)
Other co-resident earnings	-1.13 (0.81)	-0.52 (0.72)	-6.13 (2.22)	-2.50 (1.88)

Notes: Changes in household member earnings between 10 and 13 months after an initial separation for SIPP separators in the top (“most seasonal”) and bottom (“least seasonal”) deciles 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 nonpartner household adult (last row). All earnings variables are expressed as percentages of the separator’s earnings one month prior to the initial separation. All variables are adjusted on a monthly basis to ensure that the total household income change in any month is between −300% and 300%, consistent with the estimates based on winsorized total household income presented previously. Individuals with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Standard errors shown in parentheses.

Table D.4: Changes in transfers by exposure to seasonal work interruptions

	Most seasonal	Least seasonal
Household UI receipts		
<i>Unadjusted</i>	4.65 (0.40)	0.00 (0.26)
<i>Adjusted for underreporting</i>	6.93 (0.60)	-0.06 (0.37)
Other transfer receipts	0.21 (0.13)	0.10 (0.11)
SNAP receipt (p.p.)	0.56 (0.30)	0.08 (0.30)

Notes: Changes in government transfers between 10 and 13 months after an initial separation for SIPP respondents in the top (“most seasonal”) and bottom (“least seasonal”) deciles 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. Transfer income variables are adjusted on a monthly basis to ensure that the total household income change in any month is between -300% and 300% , consistent with the estimates based on winsorized total household income presented previously. Food-stamp receipt is reported in percentage points. Individuals with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Standard errors shown in parentheses.

Table D.5: Decomposing changes in household income among nonseasonal separators

	Overall	Sex of focal separator:		Type of separation:		Pre-separation sector:	
		Male	Female	Unemp.	Non-part.	Education	Other
<u>A. Unadjusted estimates</u>							
Own earnings	2.52 (0.99)	2.48 (1.27)	2.63 (1.47)	1.91 (1.25)	3.81 (1.62)	0.91 (2.88)	2.68 (1.05)
Partner earnings	1.14 (0.76)	-0.02 (0.54)	3.85 (2.19)	0.35 (0.72)	2.80 (1.82)	3.64 (3.87)	0.90 (0.74)
Other co-resident earnings	-0.52 (0.72)	-0.52 (0.86)	-0.51 (1.34)	-0.62 (0.86)	-0.29 (1.28)	-2.87 (2.94)	-0.28 (0.74)
Household UI benefits	0.00 (0.26)	0.00 (0.33)	-0.01 (0.38)	-0.05 (0.34)	0.11 (0.33)	-0.66 (0.50)	0.06 (0.28)
Other transfer receipts	0.10 (0.11)	0.04 (0.14)	0.22 (0.19)	0.12 (0.11)	0.03 (0.22)	0.07 (0.17)	0.10 (0.12)
All other sources	0.11 (0.45)	0.20 (0.49)	-0.11 (0.97)	0.48 (0.51)	-0.67 (0.88)	-0.75 (1.73)	0.19 (0.46)
Household income	3.34 (1.44)	2.18 (1.66)	6.06 (2.83)	2.19 (1.70)	5.79 (2.65)	0.34 (6.04)	3.65 (1.47)
<u>B. Adj. for UI underreporting</u>							
Household UI benefits	-0.06 (0.37)	-0.10 (0.47)	0.01 (0.57)	-0.17 (0.50)	0.16 (0.47)	-1.00 (0.78)	0.03 (0.40)
Household income	3.28 (1.44)	2.08 (1.65)	6.08 (2.84)	2.08 (1.71)	5.84 (2.64)	0.00 (6.06)	3.62 (1.46)

Notes: Changes in components of household income between 10 and 13 months after an initial separation for SIPP respondents in the bottom decile of predicted excess recurrence. All variables are expressed as percentages of the individual’s earnings one month prior to the initial separation. Monthly changes in household income are winsorized to be between -300% and 300% of pre-separation earnings, and each component of household income is adjusted proportionately on a monthly basis to ensure that they add up to the winsorized total. Individuals with pre-separation earnings less than \$450/month in January 2017 dollars are excluded. Standard errors shown in parentheses.