The Rise of In-and-Outs:
Declining Labor Force Participation of Prime Age Men

John Coglianese*

February 28, 2018

Click here for most recent version

Abstract

This paper documents that much of the decline in labor force participation of U.S. prime age men comes from "in-and-outs"—who I define as men who temporarily leave the labor force. Individuals moving in and out of the labor force have been an understudied margin of labor supply but account for roughly one third of the decline in participation between 1977 and 2015. Most in-and-outs take an occasional short break in between jobs but are otherwise attached to the labor force. Examining explanations for the rise of in-and-outs, I find that half of the rise has come from married or cohabiting men, and I show that this portion of the increase can be explained by a wealth effect from their partners' growing earnings, using variation in the growth of female wages across demographic groups. Additionally, I find that changes in household structure, particularly from young men increasingly living with their parents, account for much of the rest of the rise of in-and-outs. To examine both effects within a unified framework, I construct and estimate a dynamic model of labor supply and household formation. The model estimates imply that labor supply factors are responsible for nearly the entire rise of in-and-outs, while changes in labor demand have contributed little.

*Harvard University: coglianese@fas.harvard.edu. I am deeply indebted to my advisers Gabriel Chodorow-Reich, Lawrence Katz, and James Stock for their advice and guidance. I would also like to thank Joshua Abel, Katharine Abraham, Laura Blattner, Zachary Brown, Karen Dynan, Emmanual Farhi, Peter Ganong, Andrew Garin, Edward Glaeser, Claudia Goldin, Nathan Hendren, Xavier Jaravel, Maximilian Kasy, David Laibson, Eben Lazarus, Camille Landais, Daniel Lewis, Pascal Noel, Amanda Pallais, Mikkel Plagborg-Moller, Brendan Price, Jonathan Roth, Heather Sarsons, Jesse Schreger, Nihar Shah, Jann Spiess, Stefanie Stantcheva, Peter Tu, Chenzi Xu, and participants at the Harvard macro lunch and labor lunch for their feedback on this project.

1 Introduction

In 1960, more than 97% of American men between the ages of 25 and 54 were either employed or actively looking for work. By 2015, this rate had fallen to 88%, representing nearly 5.5 million fewer prime age workers in the labor force at any point in time. This declining trend has motivated a growing body of work seeking to understand why men are leaving the labor force.[1]

In this paper, I document that participation has changed along an understudied margin of labor supply. I find that "in-and-outs"—men who temporarily leave the labor force—represent a growing fraction of prime age men across multiple data sources and are responsible for roughly one third of the decline in the participation rate since 1977. In-and-outs take short, infrequent breaks out of the labor force in between jobs, but they are otherwise continuously attached to the labor force. Leading explanations for the growing share of permanent labor force dropouts, such as disability, do not apply to in-and-outs. Instead, reduced-form evidence and a structural model of household labor supply both indicate that the rise of in-and-outs reflects a shift in labor supply, largely due to the increasing earnings of men's partners and the growth of men living with their parents. Together, these facts paint a different picture of declining labor supply among prime age men than documented previously.

The rise of in-and-outs is not apparent from measures of the participation rate taken at a single point in time, resulting in prior work overlooking this margin of labor supply. If the participation rate falls by one percentage point, it could have been driven by one percent of men leaving the labor force permanently, or two percent spending half of their time out of the labor force, or twelve percent deciding to take one month of the year to spend out of the labor force. It is important to distinguish between these scenarios because they each may have different implications for inequality, human capital accumulation, and policy responses.

I show that the rise of in-and-outs is a robust phenomenon observed since the 1970s across a number of different longitudinal data sources. In the Current Population Survey (CPS), in-and-outs rose 6 percentage points (p.p.) between 1977 and 2015, accounting for one-third of the total decline in participation over this period. Other datasets show a similar rise in temporary non-participation over the last several decades, although the magnitudes vary somewhat across datasets. Previous studies that tried to measure in-and-outs relied on the March CPS supplement, which I show has seen increasingly inaccurate responses over the last several decades, and therefore these studies incorrectly concluded that in-and-outs were not growing as a share of the population.

In-and-outs take short, infrequent breaks in between jobs. After leaving the labor

[1] Examples include Juhn (1992; 2003), Moffitt (2012), CEA (2016), and Eberstadt (2016).

force, most in-and-outs return to the labor force within six months and are no more likely to take another break than the average man. Few in-and-outs take repeated breaks out of the labor force, meaning that most in-and-outs are highly attached to the labor force over their lifetime. While out of the labor force, in-and-outs replace time spent working with leisure activities, primarily watching television. In-and-outs do not spend much more time on child care, care for adults, educational activities, or job search during these breaks.

The rise of in-and-outs appears to be a distinct margin of labor supply from changes in permanent dropouts. The two groups cite different reasons for leaving the labor force, report different qualities of life, and show up in different regions. While some in-and-outs could turn into dropouts, the distinctions between these two groups make it more natural to treat them as separate components of the participation decline, with separate explanations, rather than intertwined aspects of the same phenomenon. Additionally, some of the most common explanations that have been put forward for the growth of permanent dropouts, such as disability insurance or incarceration, cannot explain the rise of in-and-outs.

Next, I turn to understanding the forces responsible for the rise of in-and-outs more formally. In particular, I differentiate between explanations related to changes in labor demand, reflected in diminished market opportunities, from changes in labor supply, representing lower participation holding market opportunities constant. I start by examining reduced-form evidence for supply- and demand-side explanations individually, before examining these explanations jointly through the lens of a structural model.

Consistent with a labor supply shift, I show that household structure is important for explaining the rise of in-and-outs. Using a descriptive shift-share decomposition, I find that the majority of the rise of in-and-outs stems from two sources. First, a rise of in-and-outs among married or cohabiting men is responsible for half of the total increase. Second, much of the rest of the rise of in-and-outs is accounted for by the decline in marriage rates, since men without partners are more likely to be in-and-outs. To the extent that the same forces are responsible for both of these phenomena, measuring the total contribution of these forces involves estimating both their effects on participation within households as well as their effects on reallocation across household types, motivating the later use of a structural model to estimate both of these channels simultaneously.

Focusing on the rise of in-and-outs among married and cohabiting men, I show that this rise is mostly due to a wealth effect coming from rising income of men's female partners. To provide support for this channel, I test whether this effect is strongest in households where the wages of men's partners are predicted to have risen the most, as would be implied by a wealth effect. I regress the rise of in-and-outs across demographic groups partitioned by age, race, education, and state of residence against wage growth for married and cohabiting women within these groups. To address potential endogeneity, such as reverse causality, I instrument the wages of married and cohabiting women in a group with wage growth of single women in that same group. The estimated relationship indicates that in-and-outs have risen most in households experiencing the largest growth in female wages. The magnitude of the implied wealth effect is consistent with estimates in the literature.

Furthermore, I provide indirect evidence that in-and-outs are choosing to cut back on labor supply by examining the evolution of in-and-outs' household expenditures and income. In-and-outs' household expenditures decline only slightly when leaving the labor force and then rebound when they return, suggesting that these separations do not represent unanticipated shocks to permanent income. When in-and-outs return to the labor force their labor income catches up with that of their continuously employed peers, unlike men who are fired or laid off, indicating that in-and-outs are not suffering permanent consequences of their time out of the labor force.

Turning to the role of labor demand, I find that changes in men's market opportunities can explain little of the rise of in-and-outs. Since in-and-outs are in the labor force at least some of the time, I can directly measure the market opportunities available to them and examine how they have changed.[2] While factors such as automation and offshoring have reduced market opportunities for some prime age men, average real wages have actually risen slightly since 1977, casting doubt on this explanation. Furthermore, I show that even among groups that have experienced wage declines, such as less-educated men, these declines are not nearly large enough to explain the rise of in-and-outs using conventional labor supply elasticities from the literature. In-and-outs have risen across all industries and occupations, suggesting this phenomenon is not limited to particular types of jobs. Additionally, even though changes in labor demand could directly affect employment in the short run (as is the case if wages are rigid), these effects appear to be temporary and therefore cannot explain the long-term downward trend of participation.

Finally, I examine these explanations within a unified framework to capture both the direct and indirect effects of supply and demand forces. I construct a dynamic model of labor supply and household formation. Fitting this model to observed data on participation rates, household structure, and income allows me to estimate how much of the rise of in-and-outs comes from changes in the wages available to men and women, as well as changes in factors affecting household formation. I find that the growth of female wages is responsible for half of the rise of in-and-outs, while changes in men's wages predict a slight decrease of in-and-outs. These estimates are close in magnitude to the

[2]This provides a large advantage for studying the growth of in-and-outs as opposed to dropouts, since potential wages are unobserved for the latter group. As such, it is unknown whether dropouts' potential wages have increased or decreased appreciably over the last several decades.

reduced-form results, indicating that the direct effects of wage changes on participation dominate and indirect effects are small. Most of the rest of the rise of in-and-outs is attributed to a combination of increased transfers from parents and preferences over household formation, as these forces combined account for the growth of men living alone or with their parents. In total, these findings imply that nearly all of the rise of in-and-outs reflects a decline in the desired amount of labor supply on the part of prime age men.

This paper relates to several literatures. First, the pattern of changing participation among prime age men documented in this paper contrasts with several prior examinations of the change in prime age male participation. Juhn (1992; 2003), Juhn et al. (1991; 2002), and Moffitt (2012) conclude that the decline in participation was driven almost entirely by an increase in dropouts, based on evidence from the March CPS, and they attribute this to changing market opportunities. Elsby and Shapiro (2012), Autor, Dorn and Hanson (2013), Autor and Wasserman (2013), and CEA (2016) also point to reduced market opportunities for prime age men explaining some of the decline in participation. The availability of disability insurance has been highlighted by Autor and Duggan (2003), Eberstadt (2016), and Winship (2017) as a potential explanation for the decline. Schmitt and Warner (2011) and Eberstadt (2016) also point to the rising population of men with a criminal record as a factor in lowering the share of men working. Aguiar, Bils, Charles and Hurst (2017) focus on young men and show that improvements in video game technology may account for much of the decline in participation among this group. Krueger (2017) shows that the decline in participation among prime age men may be related to the opioid crisis, as participation has fallen more in areas with higher rates of opioid prescriptions.

Documenting the rise of in-and-outs also provides evidence for a key margin of adjustment for aggregate labor supply. In-and-outs rising in response to a decrease in desired labor supply is a key prediction of the time-averaging aggregation theory of Mulligan (2001) and Ljungqvist and Sargent (2007), who introduce this approach to reconcile large aggregate labor supply elasticities with small micro elasticities. In contrast, prior theories of labor supply aggregation point more towards labor supply adjusting along the dropout margin (Hansen, 1985; Rogerson, 1988). Additionally, the rise of in-and-outs documented in this paper contributes to the discussion of homogeneity versus heterogeneity in participation rates dating back to Heckman and Willis (1977). While Goldin (1989) shows that the rapid growth of female labor supply during the 20$^{\text{th}}$ Century reflects heterogeneity in participation rates, the analysis presented in this paper shows that the decline of male labor supply in substantial part reflects homogeneity (i.e. in-and-outs).

This paper highlights the role of household channels in explaining changes in labor supply. This approach was pioneered by Becker (1965), who focuses on the importance of home production in affecting labor supply of family members.[3] While this framework has been used extensively to study female labor supply, relatively little work has applied this approach to studying male labor supply. Jones, Manuelli and McGrattan (2015) show that a calibrated unitary model of households implies that the large reduction in the gender wage gap over the last half century should have reduced male labor supply by much more than actually observed through a large wealth effect. Knowles (2012) offers a potential explanation for this puzzle through a collective model of households, as in Chiappori (1992), and shows that the data are consistent with rising female wages increasing women's bargaining power within households, resulting in no reduction in male labor supply from a wealth effect. This paper offers a complementary solution to this puzzle by showing that male labor supply has indeed fallen along certain margins due to wealth effects from rising female wages, albeit by a smaller magnitude than predicted by Jones et al. (2015).

The rise of in-and-outs also stands in contrast to a literature on declining dynamism in US labor markets. Previous studies have noted declines in job flows (Davis and Haltiwanger, 2014), geographic mobility (Molloy, Smith and Wozniak, 2011), short-term jobs (Hyatt and Spletzer, 2017), and startup rates (Haltiwanger, Jarmin and Miranda, 2013), among other measures. However, the evidence in this paper shows that an understudied dimension of fluidity—cycling in and out of the labor force—has been growing over the same time period.

The remainder of the paper proceeds as follows. Section 2 describes the decomposition of changes in participation into in-and-outs and dropouts. Section 3 outlines several key descriptive facts about in-and-outs, both about their characteristics as well as their behavior. In Section 4, I outline the role of labor supply shifts in driving the rise of in-and-outs, particularly the wealth effect from rising income of men's partners. Section 5 examines—and finds wanting—alternative explanations for the decline based on changes in labor demand, taxes, and transfer programs. Section 6 uses a dynamic model of labor supply and household formation to measure the total contributions of labor supply and demand factors to the decline. Section 7 concludes.

2 Separating In-and-Outs from Dropouts

This section describes how I distinguish between in-and-outs and permanent dropouts. I show that the prevalence of in-and-outs has more than doubled since the mid-1970s and accounts for one third of the total decline in participation since the mid-1970s. The rise of in-and-outs is robust to adjusting for previously documented sources of measurement

[3] Greenwood, Guner and Vandenbroucke (2017)provide a survey of the large literature that has followed this approach.

error and appears across a range of longitudinal datasets, including both survey and administrative sources. Furthermore, I show that the approach used by previous studies to measure temporary non-participation, relying on retrospective data from the March CPS, is increasingly biased and gives inaccurate estimates of in-and-outs. Together, these facts highlight the importance of an understudied margin of labor supply for the decline in labor force participation of prime age men.

2.1 Decomposition

I distinguish between in-and-outs, who are out of the labor force only temporarily, from dropouts, who are out of the labor force persistently. Changes along each of these margins have very different implications, making it important to distinguish between them. An increase in the number of dropouts will sharply increase income inequality, while growth of in-and-outs will have a more muted effect. Additionally, in-and-outs may be able to return to the labor force before their skills atrophy, while dropouts' skills likely deteriorate as they remain disconnected from the labor force. To the extent that being an in-and-out is a transitory state while a dropout is an absorbing state, short-term increases along each of these margins will have different long-run consequences.

I start by measuring changes in in-and-outs and dropouts using the Current Population Survey (CPS), a large nationwide survey collecting monthly labor force statistics. The CPS has four main advantages over other datasets for measuring in-and-outs. First, the CPS features a 4-8-4 rotating panel design in which individuals are interviewed for 4 consecutive months, then are not contacted for 8 months, and finally are interviewed for another 4 consecutive months. This rotating panel design enables me to examine labor force status for individuals longitudinally over eight months (covering a sixteen-month period). Second, the CPS measures labor force status monthly, which ensures that short spells of non-participation are captured. Third, results from the CPS are relatively straightforward to compare across decades due to consistency in the survey instrument and sampling procedures. Lastly, the CPS uses a substantially larger sample of households each month than other comparable surveys.[4]

Using the method developed by Drew, Flood and Warren (2014) implemented in IPUMS CPS (Flood, King, Ruggles and Warren, 2017), I match individuals' responses across interviews to construct a panel dataset. I restrict the sample to men ages 25-54 who can be matched successfully across all eight interviews, which is about 60% of all prime age men in the sample. These individuals are slightly more attached to the labor

[4] The CPS surveys about 60,000 households each month, compared to about 42,000 households in the 2008 Survey of Income and Program Participation (SIPP), approximately 10,000 families in the 2017 Panel Study of Income Dynamics (PSID), and approximately 9,000 individuals in the 1997 National Longitudinal Survey of Youth (NLSY).

force on average compared to all prime age men, but they have experienced nearly the same decline in participation as the overall population of prime age men. I label an individual as an in-and-out if he reports being in the labor force for between one and seven months out of eight, and I label an individual a dropout if he is not in the labor force in any of the eight interviews.[5]

Both in-and-outs and dropouts have grown over time. The share of the prime age male population classified as an in-and-out rose from about 6% in 1980 to more than 10% by 2015. Dropouts have also grown as a share of prime age men, rising from 3% in 1980 to more than 7% by 2015. The combination of these effects has resulted in the share of men who are always in the labor force falling by nearly 9 p.p. over this time period.

Next, I decompose the contribution of rising in-and-outs to the decline in labor force participation. I begin by writing the aggregate non-participation rate $(1 - \text{LFPR}_t)$ as a weighted average of the non-participation rates for each group, where s_g is the population share for group $g \in \{D = \text{Dropouts}, IO = \text{In-and-Outs}, AP = \text{Always Participators}\}$:

$$
\begin{aligned}
1 - \text{LFPR}_t &= \sum_g (1 - \text{LFPR}_{g,t}) s_{g,t} = 0 \cdot s_{AP,t} + (1 - \text{LFPR}_{IO,t}) \cdot s_{IO,t} + 1 \cdot s_{D,t} \\
&= (1 - \text{LFPR}_{IO,t}) \cdot s_{IO,t} + s_{D,t}
\end{aligned}
$$

Intuitively, the total number of non-participants is equal to the number of in-and-outs who are currently out of the labor force plus the number of dropouts. Taking the difference in the participation rate from one time period to the next yields an expression decomposing the participation change into the shares attributable to in-and-outs and dropouts:

$$
\Delta\text{LFPR}_t = \underbrace{((1 - \text{LFPR}_{IO,t})\, s_{IO,t} - (1 - \text{LFPR}_{IO,t-1})\, s_{IO,t-1})}_{\text{Change due to In-and-Outs}} + \underbrace{(s_{D,t} - s_{D,t-1})}_{\text{Change due to Dropouts}} \tag{1}
$$

Equation 1 enables me to separate the portion of the decline due to the rise of in-and-outs from the portion caused by dropouts.

As shown in Figure 1, the rise of in-and-outs accounts for about one third of the overall decline in participation in the CPS since 1977. The exact magnitudes of this

[5] I define in-and-outs and dropouts based on labor force participation, which includes both employment and unemployment. An alternative would be to define this groups based solely on employment. However, as I show in section 3.1, the rise of in-and-outs reflects movements between employment and non-participation rather than changing definitions of unemployment. As a result, the two approaches yield similar results. I use definitions based on labor force participation in the baseline for two reasons. First, previous research has shown that unemployment and non-participation are qualitatively different phenomena (Flinn and Heckman, 1983). Second, labor force participation is less cyclical than employment, making it easier to measure long-term trends without having to adjust for the business cycle.

Figure 1: Changes in Participation Due to In-and-Outs and Dropouts

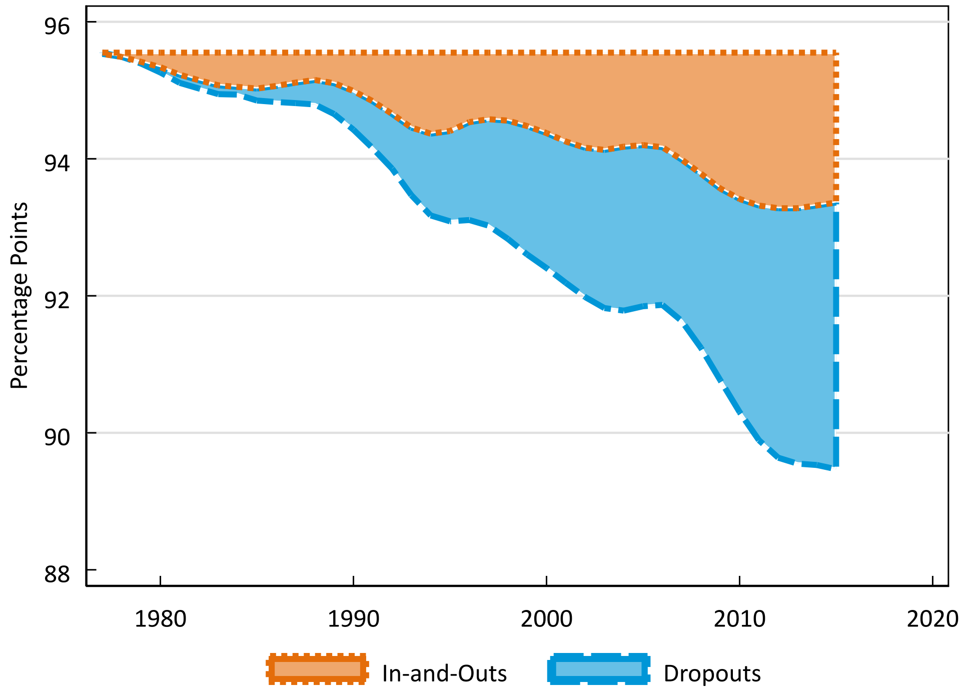

Source: IPUMS CPS matched longitudinally, 1977-2015. Graph shows the contributions of in-and-outs and dropouts to the decline in participation of prime age men. This is computed by multiplying the increase in the population share times the average non-participation rate within each group (~35% for in-and-outs, 100% for dropouts). Both series have been smoothed with an Epanechnikov kernel with bandwidth 1 to reduce year-to-year volatility and make long terms trends clearer. All statistics are computed using survey weights.

contribution are reported in Table 1, showing that in-and-outs are responsible for 2.2 p.p. of the 6.1 p.p. decline since the mid-1970s. Changes in participation due to in-and-outs have been remarkably consistent over time, declining steadily at a rate of 0.6 p.p. per decade. From the mid-1970s through the mid-1990s, the rise of in-and-outs was responsible for the majority of the decline in participation. Since then, although in-and-outs have continued to rise at the same pace, an acceleration in the growth of dropouts has accounted for relatively more of the decline in participation, especially during the aftermath of the Great Recession.

Robustness I now turn to several potential sources of measurement error, showing that in-and-outs still account for about one third of the total participation decline even after adjusting for measurement error.

One potential concern is that some of the men categorized as in-and-outs in the CPS are actually long-term non-participants, who happened to be interviewed right around

Table 1: Changes in Participation Due to In-and-Outs and Dropouts

Year	Participation Rate	Population Shares		Contributions to Participation Decline since 1977	
		In-and-Outs	Dropouts	In-and-Outs	Dropouts
1977	95.6	5.1	2.7	–	–
1980	95.3	6.0	2.7	0.3	0.1
1990	94.3	6.7	3.4	0.7	0.7
2000	92.3	8.9	4.7	1.3	2.0
2010	90.2	11.4	5.7	2.3	3.1
2015	89.6	11.1	6.6	2.2	3.9

Source: IPUMS CPS matched longitudinally, 1977-2015. All units are percentage points. Contributions to participation decline are computed using equation 1. All statistics are computed using survey weights.

the time of leaving the labor force. For example, if individuals retire in the middle of their CPS interview period they will be counted as an in-and-out since they are in the labor force for some of their CPS responses, even if they are subsequently permanent dropouts. Additionally, some of those categorized as dropouts in the CPS may be only temporarily out of the labor force, and should have been counted as in-and-outs. The true contribution of temporary non-participation can be written as

$$\Delta\text{LFPR}^{\text{Temporary}} = \alpha \cdot \Delta\text{LFPR}^{\text{CPS In-and-Outs}} + \beta \cdot \Delta\text{LFPR}^{\text{CPS Dropouts}} \tag{2}$$

where α is the share of non-participation by CPS in-and-outs that is actually temporary and β is the analogous share for CPS dropouts.

I measure α and β by recreating the CPS interview structure in another dataset where I can observe the duration of non-participation spells. I use the Survey of Income and Program Participation (SIPP), a monthly longitudinal survey of individuals. Different iterations of the SIPP have tracked individuals for varying lengths of time, with most panels typically lasting between two and five years. I define temporary non-participation as a spell lasting less than 24 months, which requires me to use data tracking individuals for 64 consecutive months.[6] Only the 2008 panel of the SIPP meets this requirement, so I can only estimate α and β for a single point in time.[7] Of those categorized as in-and-outs by the CPS definition, 81% of their non-participation is temporary as measured by the

[6] I recreate the CPS definition of in-and-outs, which aggregates participation status over a sixteen-month period (ignoring the middle eight months). To be able to determine whether non-participation is temporary, labor force status must be observable for two years before and after the CPS definition. The combination of these requirements necessitates a 64-month long panel.

[7] When I drop the requirement of a two year window on each side of the sixteen-month CPS period, I can use previous SIPP panels as well. Examining previous SIPP panels provides suggestive evidence that α and β have not changed substantially over time. However, without the window on each side of the CPS period, a substantial share of CPS in-and-outs' and dropouts' non-participation is censored at a duration of less than two years, meaning I cannot rule out substantial changes in α or β over this time period.

SIPP, as shown in Appendix Table A.1. Additionally, 6% of non-participation by CPS dropouts is temporary. In column 2 of Appendix Table A.2, I apply these rates to the contributions of in-and-outs and dropouts as measured in the CPS, using equation 2 to give an estimate of the contribution of temporary non-participation. This is nearly identical to the contribution in the baseline specification shown in column 1, indicating that potential incorrect measurement of in-and-outs in the CPS is not a substantial concern.

Columns 3-5 of Appendix Table A.2 show adjustments to the contribution of in-and-outs for additional sources of measurement error. In column 3, I show the contribution of in-and-outs where the labor force participation of CPS respondents has been "de-NUNified" to address documented problems with misclassification (Abowd and Zellner, 1985; Elsby, Hobijn and Şahin, 2015).[8] Column 4 drops any respondents whose labor force status in any of their 8 months is reported by proxy. Column 5 expands the definition of unemployment to include those individuals who are not actively searching for a job but say that they want to work, following Barnichon and Figura (2016). For all of these adjustments, I reclassify individuals as in-and-outs or dropouts based on their adjusted labor force status histories and recompute the contribution of in-and-outs. Across all of these adjustments, in-and-outs are responsible for roughly one third of the total decline in participation, as in the baseline specification.

2.2 Comparing to Other Datasets

In this section, I show how the rise of in-and-outs measured in the CPS compares to other data sources. First, I show that the rise of in-and-outs is a robust phenomenon observed across several different longitudinal data sources. Second, I find that the March CPS, which is commonly used in the literature, incorrectly shows a decline in the share of in-and-outs due to increasing recall bias over time.

2.2.1 Longitudinal Datasets

The results in the CPS can be compared to three other datasets that collect information on labor supply longitudinally. However, these datasets collect information at different frequencies, have different notions of labor supply, and use different sampling procedures. Below I describe how I measure the share of prime age men who are in-and-outs within each of these datasets.

[8]The process of "de-NUNifying" labor force participation consists of recoding flows of $N \to U \to N$ as $N \to N \to N$ and recoding $U \to N \to U$ as $U \to U \to U$.

Survey of Income and Program Participation (SIPP) The SIPP is a nationwide longitudinal survey of individuals, organized into panels lasting typically two to five years. I use monthly interviews from each panel of the SIPP conducted between 1984 and 2008, excluding the short 1989 panel. I focus on just the first twelve months of each panel both to maintain comparability across panels and to reduce the impact of attrition and I restrict the sample to 25-54 year old men who were interviewed in all twelve months. Monthly labor force participation is defined analogously to the CPS definition.[9] I label an individual as an in-and-out if he responds as being in the labor force for between one and eleven months out of twelve.[10]

Panel Study of Income Dynamics (PSID) The PSID is a longitudinal survey of families that began in 1968 and has continued to interview the same families, as well as their descendants and new family members, over several decades. I use all responses to the PSID through 2013, which includes annual responses from 1968-1997 and biennial responses afterwards. The sample includes all men ages 25-54 in the original sample as well as those who join original PSID families for the years in which they are interviewed. Labor force participation is measured at the time of the interview and is defined analogously to the CPS definition.[11] For observations before 1997, I label an individual as an in-and-out if he responded as being in the labor force for one or two of the last three years. For observations after 1997, I label an individual as an in-and-out if he responded as being in the labor force for exactly one of the last two interviews (which cover a three year period). In computing the growth of in-and-outs, I compute growth separately before and after 1997 to allow for a structural break when this definition changes.

Social Security Administration (SSA) Earnings Public-Use File The SSA Earnings Public-Use File contains annual earnings records for a 1% sample of all individuals issued Social Security numbers prior to 2007. I create measures of annual participation from SSA records of men ages 25-54 covering the 1960-2006 period.[12] I classify an indi-

[9]I denote an individual as employed in a week if he had a job or business, including if he was absent from work. I denote an individual as being unemployed if he is actively looking for work or on temporary layoff in a week. Individuals are counted as being in the labor force in a month if they are employed or unemployed for at least one week out of the month.

[10]Since I am limited to using just the first twelve months of each panel for comparability over time, the duration of non-participation is censored for many spells, preventing me from defining in-and-outs based on non-participation spell duration as I did previously. Nonetheless, the definition based on the first twelve months appears to capture almost exclusively temporary non-participation and this relationship appears to be consistent over time, as judged from those spells that are not censored.

[11]Individuals are counted as being in the labor force if they are either employed or unemployed at the time of the interview. Unemployed individuals include those on temporary layoff.

[12]While the dataset contains earnings records starting in 1951, these records are unusually volatile in the mid-1950s and produce a sharp break in the share of in-and-outs between 1958 and 1959, possibly due to changes in measurement. To avoid this break affecting the estimated growth of in-and-outs, I drop

Table 2: Rise of In-and-Outs Across Datasets

Dataset	Data Type	Time Period	Average Rate of Change (p.p. / decade)			
			In-and-Outs	Dropouts	Always Participators	LFPR
CPS	Longitudinal	1977-2015	+1.53	+1.05	-2.58	-1.67
SIPP	Longitudinal	1984-2008	+1.65	+1.21	-2.85	-1.71
PSID	Longitudinal	1968-2013	+1.54	+0.92	-2.46	-2.19
SSA	Longitudinal	1960-2006	+0.59	+1.32	-1.91	-1.52
March CPS	Retrospective	1976-2015	-1.00	+1.39	-0.40	-1.25

Note: This table shows the average rate of change in population shares for three distinct labor force categories as well as the participation rate. The average rates of change are normalized and expressed in percentage points per decade for comparability. Definitions of the labor force categories vary across data sources, as described in the text. Statistics from the CPS, SIPP, and March CPS are computed using survey weights.

vidual as being in the labor force in a given year if his annual earnings exceed half the minimum wage times 40 hours per week times 13 weeks per year. I label an individual as an in-and-out if he participated for one or two of the last three years.

Comparison All of the longitudinal datasets show rising in-and-outs among prime age men over the last three to five decades. Table 2 reports the average change in the population share of in-and-outs, dropouts, and always participators for each dataset. In-and-outs have risen as a share of the prime age male population by slightly more than 1.5 p.p. per decade in the CPS, SIPP, and PSID. In-and-outs rose at a slower average rate of about 0.6 p.p. per decade in the SSA dataset, partially due to the difference in sample period.[13] The rise in the share of in-and-outs appears fairly consistent across these datasets, despite differences in the time period and frequency at which participation is measured. Additionally, the rise of in-and-outs is observed across both survey and administrative data, suggesting that this phenomenon is not simply due to changes in survey response patterns.

2.2.2 March CPS

The March CPS Supplement is commonly used in the literature to measure labor supply. Individuals are asked to provide the total number of weeks they were employed or unemployed over the past year, which was used by Juhn et al. (1991; 2002) to distinguish between temporary and permanent non-participants. However, this information is all

records from before 1960.

[13] The SSA dataset estimates that the share of in-and-outs slightly declined between 1960 and 1970, before rising rapidly over the next several decades. If the change in in-and-outs is measured only over the 1970-2006 period, the average rate of increase is 1.2 p.p. per decade instead.

collected retrospectively. To the extent that temporary non-participation spells are less salient than long-term non-participation, they may be undercounted. Prior evidence has shown that the March CPS undercounts unemployment spells during periods when unemployment is less salient (Akerlof and Yellen, 1985).

Using the retrospective information in the March CPS, I can construct a measure of in-and-outs from this data source. Starting in 1976, the March CPS asked respondents to report the number of weeks (out of 52) that an individual was employed or looking for work in the prior calendar year. I label an individual as an in-and-out if that individual reports being in the labor force for between one and fifty-one weeks during the prior calendar year. I also use a more restricted sample limited to individuals whose March CPS responses can be matched to eight basic monthly CPS responses, which is limited to responses from 1989-2015 (more details can be found in Appendix B).

In contrast to the longitudinal datasets, the March CPS shows a decline in the share of in-and-outs over the last several decades. On average, the share of in-and-outs declined by about 1.0 p.p. per decade in the March CPS, as shown in Table 2. Relying on the March CPS alone would lead one to conclude that the entire decline in participation is attributable to dropouts, contrary to the evidence from the longitudinal datasets.

I directly test whether changes in recall errors explain the divergence between the March CPS and other datasets, finding evidence that March CPS respondents are increasingly failing to answer the survey accurately. Using the sample of March CPS responses matched to the basic monthly CPS, I compare labor force participation over the prior year as reported in the second March CPS interview to labor force participation as reported in the first four responses to the basic monthly CPS (almost all of which fall in the prior calendar year). Two types of recall errors can be unambiguously measured in the data: a) respondents who reported being in the labor force for all 52 weeks of the prior year in the March CPS, but in the basic CPS reported participating only three out of four months or fewer, and b) respondents who reported being out of the labor force for 52 weeks in the March CPS, but in the basic CPS reported participating at least one month out of four. Figure 2 displays the prevalence of these two types of errors from 1989-2015, showing that both types of errors have become more common. The increases in these unambiguous errors are large enough to explain most of the divergence between the March CPS and other datasets.

3 Descriptive Facts about In-and-Outs

Having shown that in-and-outs are rising, I now turn to describing in-and-outs' behavior and characteristics. First, I show that the rise of in-and-outs is mostly accounted for by an increase in the probability of leaving the labor force immediately after an employment

Figure 2: Recall Bias in the March CPS

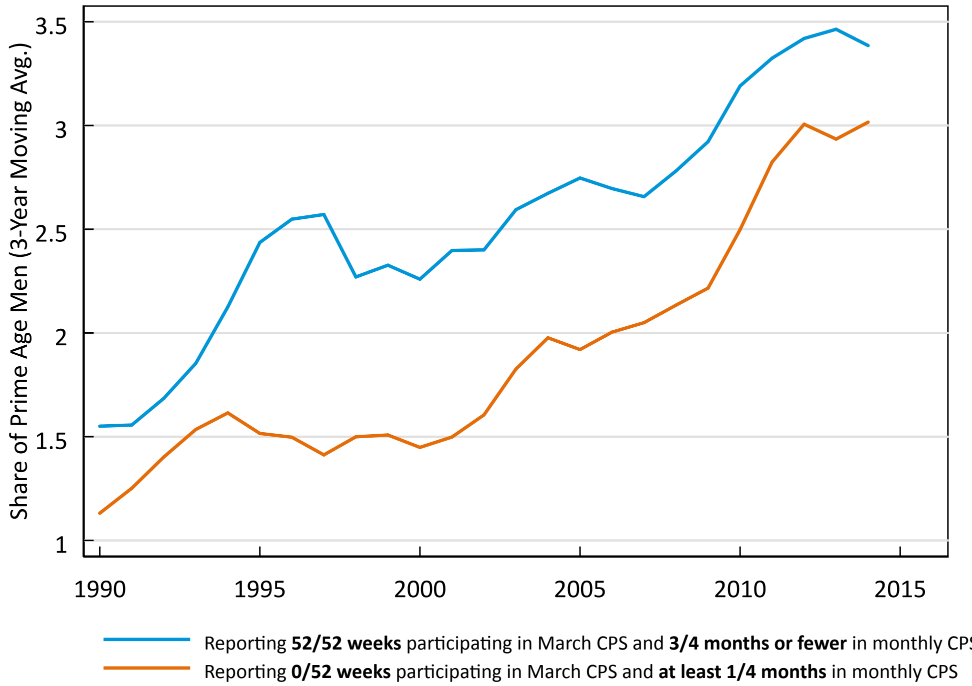

Notes: IPUMS March CPS matched to longitudinal monthly IPUMS CPS, 1990-2014. Individuals who participate in the March CPS during months 5-8 of their CPS rotation are matched to their responses during months 1-4, which all took place during the prior year. March CPS observations are matched to monthly CPS using the method of Flood and Pacas (2016). Both series are computed using survey weights.

separation, with little change in the average duration of these non-participation spells. These spells are not coming from a rise of serial non-participators, however, as once in-and-outs are back in the labor force they are no more likely to leave again than the average individual. Next, I describe the characteristics of in-and-outs and show that they appear to be qualitatively distinct from dropouts, particularly in terms of their reasons for non-participation. Lastly, I present evidence on how in-and-outs spend their time in and out of the labor force, showing that in-and-outs mostly substitute between work and leisure. Together, these facts indicate that the short, infrequent breaks of in-and-outs represent a distinct margin of labor supply from dropouts.

3.1 Changes in Labor Force Transitions

This section shows that the rise of in-and-outs reflects an increase of men taking short, infrequent breaks in between jobs. Looking at gross flows between different labor force states, the likelihood of leaving the labor force immediately after separating from a job has risen dramatically, while the rate at which men re-enter the labor force has

Figure 3: Hazard Rates of Labor Force Transitions

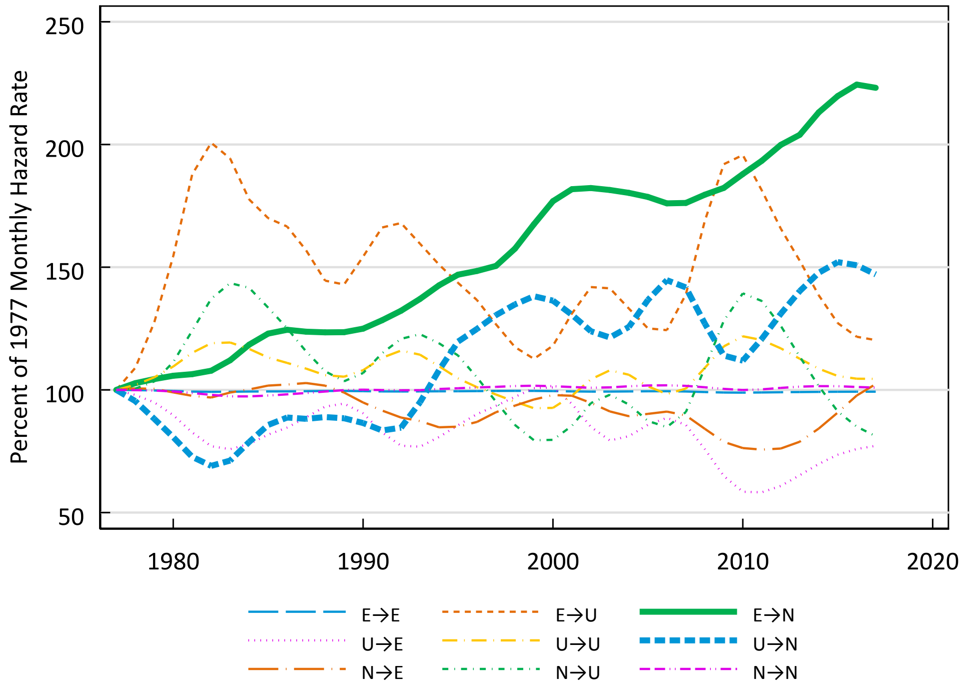

Source: IPUMS CPS, men ages 25-54 matched longitudinally for all eight responses. For each transition type, I take annual averages of the monthly hazard rates for this transition and index the series to equal 100 in 1977. All series are smoothed with an Epanechnikov kernel with bandwidth equal to 1 before indexing. All statistics are computed using survey weights.

been largely unchanged. Consistent with this finding, the average duration of non-participation spells has not changed substantially since the 1980s. Furthermore, the increase of short non-participation spells has not been due to a subgroup of men repeatedly cycling between employment and non-participation, as in-and-outs are no more likely to engage in additional non-participation spells in the future than the average individual.

Labor Force Exits I start by examining how transitions between different labor force states have changed. Using the CPS, I compute the monthly hazard rates of transitioning between employment, unemployment, and non-participation for each year, i.e. the nine transition probabilities $\pi_{ij} = \Pr\left(s_t = j | s_{t-1} = i\right)$, where $s_t \in \{E, U, N\}$. The rise of in-and-outs indicates that flows between being in the labor force and out of the labor force have increased, and examining the hazard rates can clarify which types of flows are responsible for this phenomenon.

Based on these estimated hazard rates, prime age men are increasingly likely to leave the labor force immediately following an employment separation. Figure 3 shows

Figure 4: Duration Distribution of Non-Participation

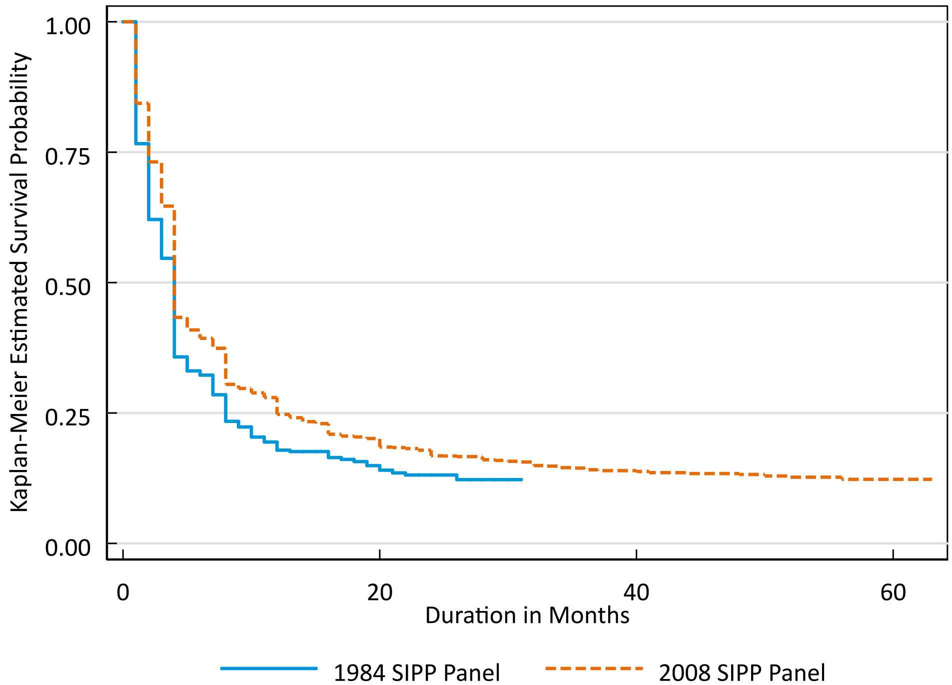

Source: SIPP men ages 25-54. Survival probabilities indicate the probability that a non-participation spell lasts at least as long as the given duration. Non-participation spells beginning in the first month an individual appears in the SIPP (left-censored spells) are excluded. All statistics are computed using survey weights.

the percentage change in these monthly hazard rates since 1977. Most of these transition probabilities have remained stable over this time period, fluctuating within 20 percent of their 1977 value. However, the probability of a man going from employment directly to non-participation has nearly doubled since 1980, rising steadily over the entire time period.[14] This indicates that the rise of in-and-outs has been primarily driven by movements between employment and non-participation directly, which reflect real variation in labor supply and not simply changing definitions of unemployment versus out of the labor force.

Non-Participation Spell Durations The amount of time that in-and-outs spend out of the labor force has not changed substantially in the last few decades. I measure the duration that in-and-outs spend out of the labor force using samples of prime age men

[14]While the probability of unemployment to non-participation transitions has risen similarly, the number of men who are unemployed is much smaller than the number who are employed, meaning that the increase in this probability is less consequential than the increase in the employment to non-participation transition rate.

from the 1984 and 2008 SIPP panels.[15] For each panel, Figure 4 shows the estimated Kaplan-Meier survival curve for non-participation spells, which measures the fraction of in-and-outs who have returned to the labor force within a given duration.[16] Most prime age men who leave the labor force return within 4 months, with 75% returning within one year. The duration of non-participation spells has slightly increased between 1984 and 2008, mainly for spells between 8 and 24 months long, but the two distributions are quite close. Little of the rise of in-and-outs can be attributed to a change in the duration of non-participation among in-and-outs; instead, the rise is mainly due to more men moving from employment to non-participation in the first place.

Persistence of In-and-Outs Once in-and-outs have returned to the labor force, they are no more likely to leave again than average. To examine in-and-outs' long-term labor supply, I use a sample of prime age men from the PSID, which allows me to examine labor force participation over many years for each individual.[17] I compute for each individual the total number of temporary non-participation spells lasting two years or less between ages 25 and 54. Figure 5 plots the distribution of these in-and-out spells across individuals. Slightly more than 20% of all prime age men experience at least one of these spells during their prime working years, with about 5% experiencing two or more such spells. However, among men who have been an in-and-out at least once, the distribution of additional spells is nearly identical to the unconditional distribution. As a result, the conditional distribution shown in Figure 5 is almost exactly equal to the unconditional distribution shifted by one to the right. This indicates that most in-and-outs are not serial non-participators; once back in the labor force they are no more likely to experience another non-participation spell than the average individual.

3.2 Who are In-and-Outs?

This section lays out the characteristics of in-and-outs compared to other groups. In-and-outs are lower-skilled on average than always participators, but higher-skilled than dropouts. In contrast to dropouts, few in-and-outs report disability-related reasons for

[15] I use the SIPP instead of the CPS so that I can estimate the duration distribution non-parametrically, which requires long uninterrupted panels of monthly labor force status.

[16] A concern with survival curve estimates is the potential for missing data. While these panels last for several years each, some individuals may not have returned to the labor force by the end of the survey, resulting in non-participation spells where duration is right-censored. Other individuals begin the panels as non-participants, giving rise to left-censored spells. I include right-censored spells in estimating the survival curves and adjust for their censoring. I exclude left-censored spells since the survival curves cannot be adjusted for censoring in these cases without imposing parametric assumptions.

[17] In the CPS I am limited to observing individuals' labor force participation only over a one-and-a-half year period and in the SIPP I can observe labor force participation for up to five years. In contrast, the PSID measures labor force participation over multiple decades for many respondents.

Figure 5: Distribution of Lifetime In-and-Out Spells

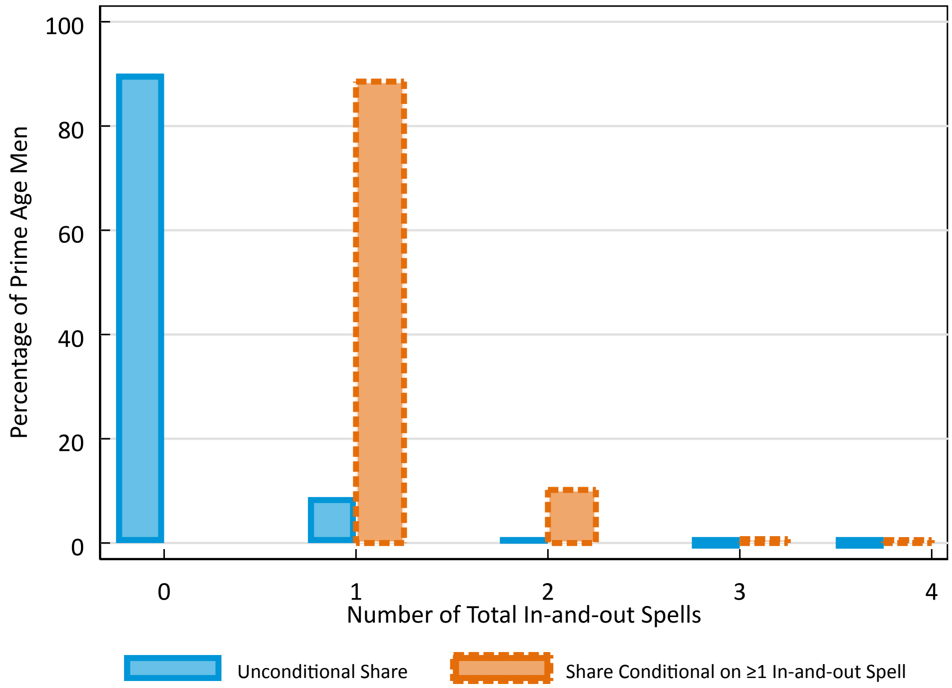

Source: PSID, men ages 25-54 with at least 10 years in sample. In-and-out spells are defined as an employment separation followed by a period of non-participation of two years or less. Before 1997, this means an individual is employed in period $t-1$, non-participating in t, and participating again in $t+1$ or $t+2$. After 1997, this means an individual is employed in period $t-2$, non-participating in t, and participating again in $t+2$. For each individual, I compute the total number of such spells they experience between ages 25 and 54. From this, I compute the distribution of the number of spells across individuals.

being out of the labor force, few take pain medication daily, and few live in economically declining regions.

Demographics I start by comparing the characteristics of dropouts, in-and-outs, and always participators. Table 3 shows average demographic characteristics for each of the three participation categories pooling over all years in the CPS. On average, in-and-outs are somewhat lower skilled (only 45% with at least some college) compared to always participators (58%), but higher skilled than dropouts (32%). Dropouts are about three years older on average than the other two categories. Both in-and-outs and dropouts are slightly more likely to be black or Hispanic than always participators.

Several important patterns emerge with regard to household structure. Nearly half of dropouts and more than half of in-and-outs are living with a married or cohabiting partner, whereas the rate for always participators is substantially higher (79%). However, dropouts and in-and-outs are much more likely to be living with their parents than

Table 3: Demographics of In-and-Outs, Dropouts, and Always Participators

	Dropouts	In-and-Outs	Always Participators
Age	42	39	39
% Black	18	17	8
% Hispanic	10	15	10
% College	32	45	58
% Married or Cohabiting	46	56	79
% Household Head without Partner	23	16	13
% Living with Parents	27	22	6
# of Children	0.7	1.0	1.3
Family Size	3.0	3.4	3.4
Household Income (2016$)	44,061	64,972	96,940

Source: IPUMS CPS matched longitudinally, 1977-2015. March CPS 1987-2015. Black share includes individuals who are both black and Hispanic, while Hispanic share does not. College includes individuals with some college but no four-year degree. Married or cohabiting individuals either report that they are a spouse or unmarried partner, or another individual within the household reports that they are the spouse or unmarried partner of the individual in question. Household heads without partners are those who report being the head of their household but are not married or cohabiting. Living with parents denotes that individuals are the child or other younger relative of the household head. Household income is measured from the IPUMS March CPS and matched to the longitudinal monthly IPUMS CPS using the method of Flood and Pacas (2016), after being deflated with the PCE price index. All statistics are computed using survey weights.

always participators. Only about a sixth of in-and-outs are the sole head of their household, which implies that the vast majority of in-and-outs are living in households with multiple potential earners.

Reasons for Non-Participation In-and-outs and dropouts report being out of the labor force for very different reasons. In the CPS, non-participants are asked about their main activity and their reasons for non-participation. I combine these into six mutually exclusive categories and compute the shares of in-and-outs and dropouts who cite each reason, as shown in Table 4. Relatively few in-and-outs report being retired, disabled, or ill, and these shares have not changed much in the last few decades. Instead, the rise of in-and-outs has all come from non-participants who say they are either in school, taking care of their house or family, or out of the labor force for some other reason.[18] In contrast, the majority of dropouts report that they are disabled and this category accounts for most of the growth of dropouts over the last few decades.

[18]Despite the growth of in-and-outs citing school as the reason for non-participation, the share of in-and-outs who are actively attending school is very small and has not risen substantially. I measure attendance in any type of schooling from the CPS October Education Supplement.

Table 4: Reasons for Non-Participation

Self-Reported Reason	Dropouts	In-and-Outs
Disability	74.6	24.3
Sickness	0.6	3.3
Retired	10.4	5.6
Taking Care of House/Family	6.6	21.8
In School	4.9	13.9
Other	2.9	30.9

Source: IPUMS CPS matched longitudinally, 1993-2015. For both dropouts and in-and-outs, I compute the share of individuals within that category who report each reason for non-participation (i.e. shares add to 100% for both dropouts and in-and-outs). Shares are pooled over the whole time period. Non-participating individuals are first asked if they are disabled, retired, or other, and, if the latter category, they are asked what their main activity during the reference week was. The reason for non-participation is only available for in-and-outs during months in which they are out of the labor force, overweighting in-and-outs that are out more often, so I reweight these observations to count all in-and-outs equally. All statistics are computed using survey weights.

Quality of Life In-and-outs report having a higher quality of life than dropouts. Appendix Table A.3 shows several measures of quality of life from the American Time Use Survey (ATUS), which has been matched to the monthly CPS to separate in-and-outs from other groups. In-and-outs report a higher level of life satisfaction on a 10-point scale compared to dropouts, averaging 6.3/10 and 5.9/10 respectively.[19] Dropouts also appear to have higher levels of pain and disabilities. A majority of dropouts (61%) report taking some form of pain medication yesterday, compared to 29% for in-and-outs and only 19% for always participators. Dropouts are also much more likely than in-and-outs or always participators to report experiencing physical or cognitive difficulties.

Geographic Distribution In-and-outs and dropouts are regionally distinct. Appendix Figure A.1 shows the change in participation along each of these margins for all 50 states and DC over the period 1980-2013. The rise of in-and-outs has not correlated with growth of dropouts across states. Dropouts have grown most in states like Kentucky and West Virginia as well as many states in the Rust Belt. Meanwhile, in-and-outs have grown across many different regions of the US, with the largest increases of in-and-outs coming in New Mexico, Alabama, Delaware, and New York. Dropouts appear to be rising most in states that have seen declining opportunities relative to other states over this time period, but no such clear pattern emerges for in-and-outs.

[19]The 10-point Cantril ladder measure used by the ATUS is described in more detail in Krueger (2017).

Table 5: Time Use of In-and-Outs While In and Out of the Labor Force

Activity	Hours per Weekday		Activity	Hours per Weekday	
	In LF	Out of LF		In LF	Out of LF
Child Care	0.3	0.5	Leisure	4.1	7.2
Care for Adults	0.1	0.1	*Watching TV*	*2.7*	*4.5*
Education	0.3	0.5	*Computer Use*	*0.2*	*0.4*
Household Activities	1.1	2.1	*Video Games*	*0.1*	*0.2*
Personal Care	8.7	10.0	*Socializing*	*0.5*	*0.7*
Health-related care	*0.0*	*0.3*	Job Search	0.1	0.1
Sleeping	*8.0*	*9.3*	Working	6.6	0.1

Source: IPUMS American Time Use Survey, 2003-2015, matched to basic monthly IPUMS CPS.
Note: Subcategories listed in grey italics. Categories are not exhaustive, so time use may not add to 24 hours. Unemployed in-and-outs have been excluded for comparability. Unemployed in-and-outs have time use similar to men out of the labor force. Each category includes travel time associated with that activity. All statistics are computed using survey weights.

3.3 In-and-Outs' Time Use

Examining how in-and-outs spend their time out of the labor force can provide insight about their lives. When in-and-outs leave the labor force, they might spend their increased non-work time on child care, care for older adults, or increased home production. Alternatively, time previously spent on work could be spent on leisure activities. I construct measures of time spent on different activities from the American Time Use Survey (ATUS) through IPUMS (Hofferth, Flood and Sobek, 2017). Since the ATUS is a follow-up survey from the CPS, I can match ATUS responses to individuals' labor force participation during their 8 months in the CPS. This allows me to examine the time use of in-and-outs who are subsequently in the labor force versus out of the labor force at the time of the ATUS interview.

In-and-outs replace time spent on work mostly with additional leisure activities. When working, in-and-outs work an average of 6.6 hours per weekday. Upon leaving the labor force, this time is mainly split between leisure activities (increase of 3.1 hours), sleeping (1.3 hours), and household activities (1.0 hour). In-and-outs do not substantially increase their time spent on child care, care for adults, education, or health-related care. Out of the three-hour increase in leisure, more than two thirds is spent watching TV, bringing total TV watching up to nearly 5 hours per day on average. In-and-outs only slightly increase their time spent on video games when leaving the labor force.[20]

[20]This finding is in contrast to recent results by Aguiar et al. (2017), which suggest that declining participation can accompany increasing time spent playing video games. In Appendix D, I decompose the contributions of different restrictions to the divergence between these results, finding that the main reasons for the difference in results are that I focus on in-and-outs as opposed to all non-employed individuals and that I examine a sample of men ages 25-54 instead of men ages 21-30.

4 Changes in Labor Supply

This section shows that most of the rise of in-and-outs can be attributed to a shift in labor supply among prime age men. I begin by arguing that the growth of in-and-outs can be modeled as a response to supply and demand factors, as in a classical framework. I then use a descriptive decomposition to show that the presence of other potential earners within the household is important for explaining the rise of in-and-outs, suggesting a labor supply shift. Half of the rise of in-and-outs comes from married or cohabiting men and most of the remainder comes from an increase of men living with their parents. I show that the former phenomenon can be explained by a wealth effect from rising earnings of men's female partners by using variation in the growth of female wages across demographic groups. Lastly, as indirect evidence for the role of labor supply, I show that the evolution of in-and-outs' household expenditures and labor income is consistent with these men taking voluntary breaks.

4.1 Modeling In-and-Outs' Labor Supply

Before investigating the evidence for a labor supply shift, it is worth asking why a labor supply shift would lead to a rise of in-and-outs rather than changes along some other margin. In a classical labor supply framework, a representative agent chooses total hours of work for the population as a function of wages and unearned income.[21] Since this setup is agnostic about the underlying margins along which labor supply changes, it is not clear which aspects of labor supply will change in response to a wealth effect or other labor supply shift. The facts laid out in the previous sections point to a reduction in labor supply along a very specific margin—an increase of short, infrequent breaks out of the labor force. It is worth considering why a labor supply shift could show up along this precise margin without changes along other margins.

If individuals smooth leisure over time, we might expect a labor supply shift to reduce hours worked per week.[22] However, if individuals face fixed costs of working during a week, e.g., due to time spent commuting, this may not be the margin that adjusts when desired labor supply changes. Prescott, Rogerson and Wallenius (2009) and Rogerson and Wallenius (2013) show that more generally if the productivity of hours worked per period features a non-convexity, as is the case if there are fixed costs, the intensive margin of labor supply is fixed by the parameters governing the non-convexity and instead the extensive margin of labor supply adjusts to equate the marginal rate of substitution with the after-tax wage. This suggests that we should expect a change in

[21] Notable examples using this framework include Lucas (1977) and Kydland and Prescott (1982).

[22] Indeed, this is in line with the famous prediction of Keynes (1930), who expected future generations to work only fifteen hours per week, although he did not mention the number of weeks worked per year.

desired labor supply to result in more weeks spent on leisure, with no change in hours per week.[23]

One might expect a reduction in desired labor supply to result in a few weeks spent on leisure every year, yet in-and-outs' breaks appear to be infrequent, with few in-and-outs repeatedly taking time out of the labor force. Part of the explanation may lie in the fact that in-and-outs appear to be highly attached to the labor force while working; most hold jobs for many years at a time.[24] In-and-outs' breaks may be an opportunity for men to consider alternative career paths or recharge before their next job, in which case they would occur at most only every few years. As a result, in any given year some men are taking a break out of the labor force even though few are doing this repeatedly.

The combination of these factors suggests that a labor supply shift could lead to growing in-and-outs, without any changes in labor supply along other margins. More generally, this indicates that traditional determinants of labor supply, such as the returns to work and the level of unearned income, affect the number of in-and-outs. Therefore, I consider standard forces affecting labor supply and, in the next section, labor demand to explain the rise of in-and-outs.

4.2 The Role of Household Structure

One factor that could affect a labor supply shift is the presence of other potential workers within the household. This section uses a descriptive decomposition to demonstrate that household structure is important for explaining the rise of in-and-outs, providing suggestive evidence for this type of labor supply shift. I conduct a standard shift-share decomposition across household types to capture both changes in participation within household types as well as changes accounted for by reallocation across household types.

I start by dividing men by household structure. I create four mutually exclusive categories: 1) men with a partner within the household, 2) heads of un-partnered households, 3) men living with their parents, and 4) all other men.[25] In 1990, the first group

[23]Consistent with this, hours worked per week among employed prime age men have changed little over the past several decades, rising by only 0.07% between 1980 and 2015.

[24]Using the 1984 SIPP panel, I examine the job tenure of currently employed workers who subsequently leave the labor force. This group, primarily composed of in-and-outs, has an average job tenure of about five years. The average tenure for workers who are continuously in the labor force is higher at about eight years, although much of this difference is due to the long tail of workers with tenure in excess of ten years.

[25]I have explored further divisions, including separate categories for heads and spouses, splitting on marital status, or using the presence of children in the household, but these divisions are not necessary to capture the key first-order trends. Therefore, for simplicity, I present results here based on the split into the four categories described above.

Table 6: Decomposition of Participation by Household Types, 1977-2015

	Reallocation	Within
Heads with Partners	0.0	-1.2
Heads without Partners	-0.2	-0.1
Living with Parents	-0.5	-0.1
Other	-0.0	-0.0
Total	-0.7	-1.4

Source: IPUMS CPS, 1977-2015. Individuals are assigned to household type based on the modal household type reported over their 8 months in the sample. The reallocation and within-group components are computed as described in equation 3. All statistics are computed using survey weights.

comprised 76% of prime age men, the second comprised 13%, the third comprised 8%, and the last comprised less than 3%.

The change in the participation rate due to in-and-outs can be decomposed into changes within household types and reallocation across household types. Letting g index household types and t index time periods with s_{gt} as the share of men in household type g in year t, and LFP_{gt} as the participation rate within group g in year t excluding dropouts (so that changes in this rate reflect only changes in in-and-outs), the change between any two years can be written as

$$
\begin{aligned}
\Delta LFP_{t+j} &= \sum_g s_{gt+j} LFP_{gt+j} - \sum_g s_{gt} LFP_{gt} \\
&= \sum_g (s_{gt+j} - s_{gt}) LFP_{gt+j} + \sum_g s_{gt}(LFP_{gt+j} - LFP_{gt}) \\
&= \sum_g \underbrace{(s_{gt+j} - s_{gt})(LFP_{gt+j} - LFP_{t+j})}_{\text{Reallocation Component}} + \sum_g \underbrace{s_{gt}(LFP_{gt+j} - LFP_{gt})}_{\text{Within Component}} \qquad (3)
\end{aligned}
$$

where the second line adds and subtracts by $\sum_g s_{gt} LFP_{gt+j}$ and rearranges. The third line adds in $\sum_g (s_{gt+j} - s_{gt}) LFP_{t+j}$, using the fact that $\sum_g (s_{gt+j} - s_{gt}) = 0$ and that the aggregate participation rate LFP_{t+j} is the same for all groups. Each group contributes to the change in the participation rate through two components. The term labeled "Reallocation Component" in equation 3 represents the change in participation due to men moving into or out of groups with participation rates different than the average participation rate. The "Within Component" represents the change due to participation rates changing within a particular household type, holding the population share of that household type constant.

Both within-group changes and reallocation across household types account for the rise of in-and-outs. Table 6 shows the contribution of each component by household type to the change in participation due to in-and-outs over 1977-2015. The within-group change among men with partners accounts for about half (1.2 p.p.) of the total decline

Figure 6: Wages of Prime Age Men and Women

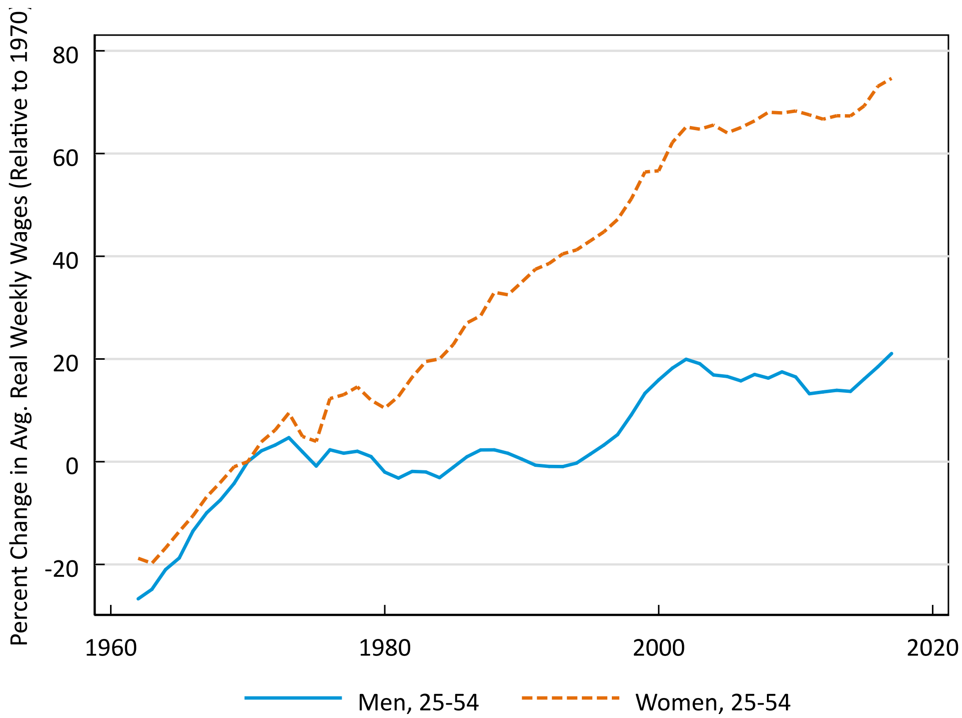

Source: IPUMS March CPS, 1962-2017. Nominal wages are deflated by the PCE price index. Weekly wages are computed by dividing annual wage income by weeks worked for those individuals who worked at least some of the prior year. For years before 1976, when weeks worked is only reported in intervals, I impute an estimate of weeks worked based on the average value within the interval among post-1976 observations. Both series are computed using survey weights.

in participation due to in-and-outs (2.2 p.p.). The reallocation of men across different household types accounts for an additional third of the rise of in-and-outs. Of the total 2.2 p.p. decline in participation, 0.7 p.p. is attributed to the reallocation component, with most of this (0.5 p.p.) coming from the increase of men living with their parents. In total, the within-group change among men with partners and the reallocation of men across household types put together account for nearly all of the rise of in-and-outs.

4.3 Wealth Effects of Rising Partner Income

Wages of mens' female partners rose dramatically over the latter half of the 20th century. While men's wages initially kept pace with women's wages, from 1970 onwards women's wages grew substantially faster, as shown in Figure 6. As women's wages rose and women joined the workforce in greater numbers, household income in dual-earner households rose accordingly. Among these dual-earner households, one possible effect of this rising wealth could be that it induced men to work less and take time off from the labor force periodically.

I test the role of wealth effects using variation across demographic groups. Using a sample of households containing a prime age man with either a spouse or unmarried female partner from the CPS, I divide households into groups based on state of residence and the age, race, and education level of the man.[26] Within each group g and for each year t, I compute the average male participation rate (excluding dropouts), $n_{m,g,t}$, and the average female post-tax hourly real wage, $w_{f,g,t}$.[27] I then regress male labor supply on female wages controlling for group-level fixed effects:

$$n_{m,g,t} = \beta w_{f,g,t} + \gamma_g + \epsilon_{g,t}$$

With group-level fixed effects, a negative coefficient β indicates that participation fell by more in groups where women's wages rose by more.

To identify the equilibrium response, I need to account for the potential endogeneity of women's wages in the presence of shocks to men's market opportunities. If men's labor supply receives a negative shock, then one possible response could be women choosing to work in higher paying but less enjoyable jobs, which could produce a negative estimate of β even in the absence of any wealth effects. This is an important concern since I am using a large number of relatively fine groups. To obviate this concern, I instrument for the wages of married women using the average post-tax hourly real wages of single women in the same group, $w^{Single}_{f,g,t}$.[28] The validity of this instrument depends on the exclusion restriction that the wages of single women are not directly correlated with the labor supply of married men, which holds if the variation in the growth of single women's wages comes from reductions in discriminatory forces across different demographic groups as in Hsieh, Hurst, Jones and Klenow (2013).[29] The first stage equation

[26]Three categories for age are used (25-34, 35-44, and 45-54), four categories for race (non-Hispanic white, black, non-black Hispanic, and all others), four categories for education (less than high school, high school graduate, some college, and college graduates), and 51 for states (including DC). This gives a total of 2,448 possible categories, although in any given year some of these cells are empty.

[27]The participation rate is calculated excluding dropouts, so that variation only comes from changes in in-and-outs. Post-tax wages are computed from multiplying the average CPS ORG wage within a group by the average net-of-tax rate in the group computed from March CPS data and the NBER TAXSIM calculator.

[28]For a married woman in a demographic group g', I compute the average post-tax real hourly wages of single women in the same demographic group, $w^{Single}_{f,g',t}$. I then aggregate these wages to the demographic group of the husband, g, by averaging across all married women married to husbands in demographic group g, yielding $w^{Single}_{f,g,t}$, which represents the average wages of single women with the same demographics as women married to men in group g.

[29]One potential concern is that the variation in the growth of single women's wages could in part reflect broader labor demand shocks, which would be correlated with the wages of married men and thus violate the exclusion restriction. I can examine this concern directly by repeating this regression with married men's wages as the outcome. I find a weakly positive effect on married men's wages, suggesting that in-and-outs should have, if anything, decreased slightly among this group absent a wealth effect.

Figure 7: Declining Participation due to In-and-Outs vs. Change in Partner's Wage

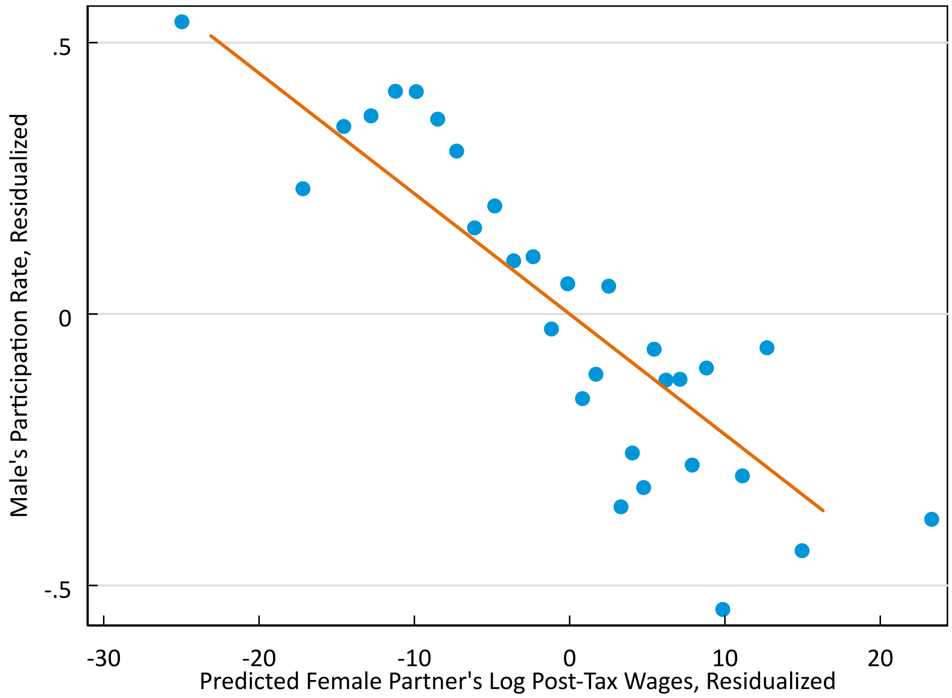

Source: IPUMS CPS, 1989-2015. Prime age men with partners are divided into 2,448 groups based on education, race, age, and state. Dropouts are excluded from this analysis. Blue dots show a scatterplot of the data broken into 30 equal size bins, plotting the mean of each variable within each bin. Female partner wages are predicted from wages of single women in the same group, controlling for group fixed effects. Data are residualized on group fixed effects before binning and plotting. Estimated coefficient equal to -0.0228 (0.0019). All statistics are computed using survey weights.

is given by:

$$w_{f,g,t} = \alpha w_{f,g,t}^{Single} + \delta_g + \nu_{g,t}$$

By using the wages of single women to instrument for the wages of married/cohabiting women, the coefficient β can be interpreted as measuring the strength of the wealth effect on male labor supply.

Across households, a clear pattern emerges of in-and-outs rising in households where women's wages increased more. Figure 7 shows a binned scatterplot of the second stage of the regression, where the observations are divided into 30 equal size bins and within each bin $n_{m,g,t}$ is plotted against the predicted values from the first stage $\hat{w}_{f,g,t}$. The relationship is summarized by the solid orange best-fit line, which indicates that the equilibrium response is downward-sloping. Households with larger predicted increases in female wages also experienced a larger rise of in-and-outs. The slope indicates that a 10% predicted increase in female wages reduced the male participation rate by 0.2 p.p. through the in-and-out margin.

Figure 8: Predicted Rise of In-and-Outs from Wealth Effect

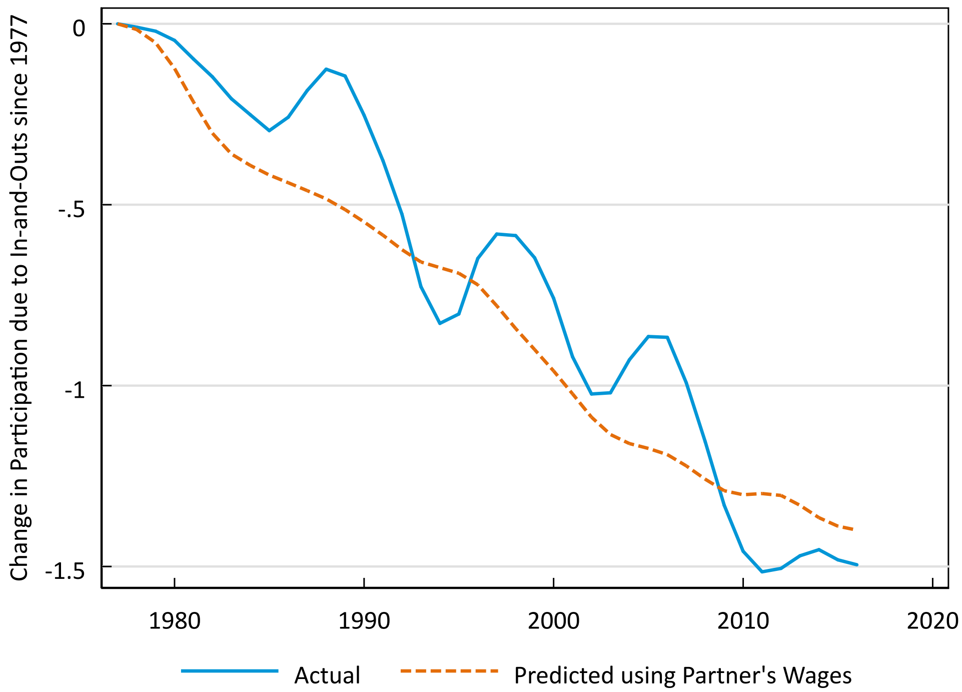

Sample: 25-54 year old men in the CPS with a spouse or unmarried partner, excluding dropouts. The predicted change in participation is computed from the change in spousal post-tax real hourly wages times -0.0228. Both series are smoothed with an Epanechnikov kernel with bandwidth equal to 1. Participation and wage series are computed using survey weights.

The magnitude of this effect is in line with the predicted wealth effect using estimates from the literature. In Appendix C, I outline a simple model of labor supply in a married household to show that the response of male labor supply to changes in female wages depends primarily on the male income elasticity of labor supply and the female uncompensated elasticity of labor supply. Using estimates of these parameters from the literature, I find that the slope of the predicted wealth effect is approximately -0.02, which is very close to the estimated slope across demographic groups.

Furthermore, the estimated wealth effect accounts for the entire rise of in-and-outs among men with partners, which accounts for half of the overall decline due to in-and-outs. Using the estimated strength of the wealth effect from above, I compute the predicted rise of in-and-outs among men with partners using the observed increase in the average wage of their partners. Figure 8 shows this predicted decline in participation due to in-and-outs with the orange dashed line, which captures nearly all of the decline in participation due to in-and-outs from 1977 to 2015. This result suggests that the wealth effect from the rising wages of men's partners can explain the entire rise of in-and-outs among men with partners.

This evidence of wealth effects contrasts sharply with the conclusions of prior work. In particular, the analysis of Juhn et al. (1991) suggests that participation fell more in households where women's earnings rose the least. There are two novel elements of the analysis presented here that lead to the opposite conclusions from prior work. First, I separate changes in participation due to in-and-outs from changes due to dropouts. This is important because dropouts are both less likely to have a partner at all compared to in-and-outs and are more concentrated in the lowest skill categories, where women's earnings have not risen as much. The second novel aspect of the analysis presented here is the use of post-tax wages as the measure of returns to work. While pre-tax wages and earnings have risen the most for the highest-skilled women, tax rates for higher-skilled households have changed little while households lower in the income distribution have experienced substantial tax cuts, resulting in middle-skill households experiencing the largest rise in post-tax female earnings. By focusing on in-and-outs and using post-tax wages to measure the returns to work, I find evidence in support of a wealth effect that prior work had not uncovered.

4.4 Evolution of Household Expenditures

Examining in-and-outs' expenditures as they leave and re-enter the labor force can provide some indication about whether these transitions are insured or not. If in-and-outs are experiencing uninsured and unanticipated shocks to their long-term market opportunities that depress their permanent income, expenditures may decline when in-and-outs leave the labor force. On the other hand, if in-and-outs market opportunities are constant and they are choosing to leave the labor force, expenditures are more likely to be unchanged.

To measure the evolution of expenditures for in-and-outs, I use data from the PSID, which contains longitudinal information on household expenditures on food and individual labor force participation.[30] Food expenditure is used here as a proxy for household consumption.[31] To measure the change in expenditures of in-and-outs relative to continuously employed individuals, I regress the k-period change in log household food expenditures on an indicator for leaving the labor force, repeating this over different horizons k:

$$\Delta \log \left(Exp\right)_{i,t,t+k} = \beta^{(k)} \mathbf{1}\left(\text{In-and-Out}_{i,t}\right) + \delta X_{i,t} + \Delta\epsilon_{i,t,t+k} \tag{4}$$

where $X_{i,t}$ is a vector of individual controls containing a cubic in age as well as time

[30] Food expenditure in the PSID includes both purchases of food for consumption at home and spending on food away from home.

[31] I have applied similar analyses to data on total household expenditures in the PSID (post-1997 data only) and the CEX Interview Survey, although this results in a substantial reduction in sample size. The results are qualitatively similar, but somewhat less precise.

period fixed effects. The dummy variable $\mathbf{1}(\text{In-and-Out}_{i,t})$ is equal to 1 if the individual i is employed in year t, out of the labor force in year $t+1$, and back in the labor force in either year $t+2$ or $t+3$. I vary k between -2 and 3, giving estimates of the evolution of expenditures in a five-year window around the labor force exit.

I also compute the evolution of expenditures for two different comparison groups. First, individuals who become unemployed involuntarily can provide a benchmark estimate of consumption changes for individuals experiencing large shocks to permanent income. An extensive literature in public finance has documented that household expenditures fall substantially when individuals become unemployed, consistent with a change in permanent income.[32] I denote a man as an unemployed job loser if he was fired or laid off, or if his firm exited. Second, retirements provide a benchmark for voluntary reductions in labor supply. Retirees decrease spending slightly when leaving the labor force, mostly due to work-related expenses, but measures of consumption stay unchanged (Aguiar and Hurst, 2005). I denote a man as a retiree if he is between the ages of 62 and 68 and is employed in year t and out of the labor force in years $t+1$ through $t+4$.

In-and-outs' expenditures fall only slightly when leaving the labor force and rebound when they rejoin the labor force. Figure 9 shows the estimated evolution of expenditures for in-and-outs as well as unemployed job losers and retirees. Household expenditures are evolving similarly for in-and-outs and other men in the two years before the in-and-out leaves employment. In the first year following a separation, household expenditures of in-and-outs drop by about 5% on average, before subsequently rebounding around the time that in-and-outs have returned to work. On average, household expenditures almost fully recover within three years of an in-and-out leaving the labor force.

This strongly contrasts with the consumption behavior of workers who become involuntarily unemployed, who not only experience a larger initial consumption decline but also see no recovery in their consumption for years afterwards. However, the initial drop of in-and-outs' expenditures closely matches the drop in spending at retirement. The consumption patterns documented here are consistent with in-and-outs' permanent income remaining mostly unchanged when they exit the labor force, with a small drop in consumption due to fewer work-related expenses that rebounds when in-and-outs return to work.

[32] Studies documenting this fact include Gruber (1997), Browning and Crossley (2001), Chetty and Szeidl (2007), Hendren (2017), and Ganong and Noel (2015).

Figure 9: Consumption of In-and-outs

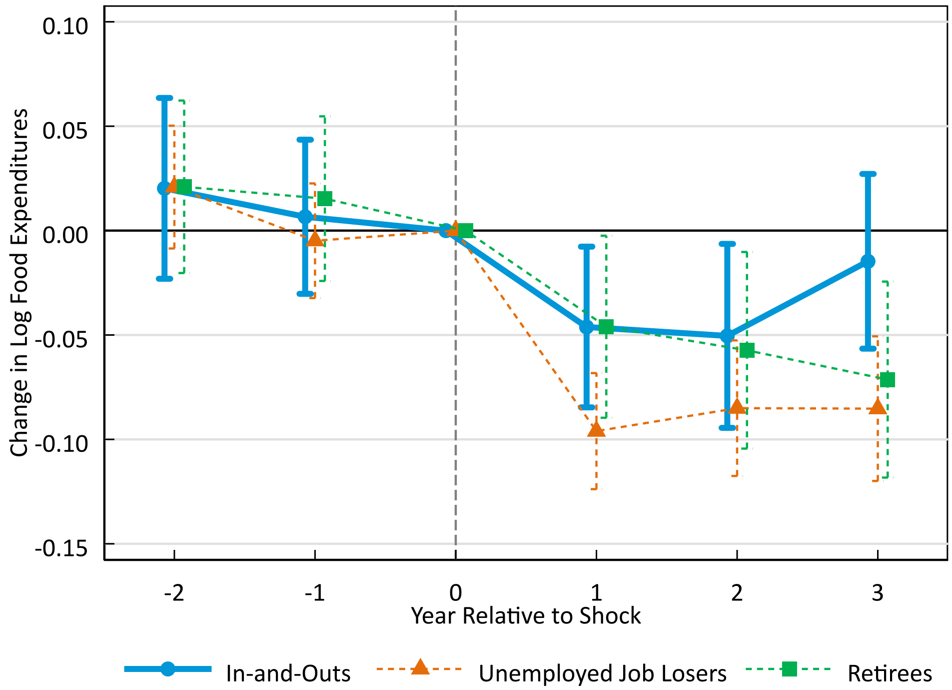

Sample: Men in the PSID from 1968-1997. In-and-outs are prime age men employed in year 0, nonparticipating in year 1, and participating again by at least year 2 or 3. Unemployed are prime age men employed in year 0 and unemployed job losers in year 1. Retirees are men ages 62-68 employed in year 0 and non-participants after. For each group, I plot the coefficients from estimating equation 4 over $k = -2, \ldots, 3$. Since I control for year fixed effects and a cubic in age, the change in log food expenditures is relative to continuously employed individuals of the same age in the same year. Standard errors are clustered at the individual level. Food expenditure includes food purchased for consumption at home as well as food away from home. I exclude observations with an annual change in log food consumption in excess of 3 to avoid problematic measurement error.

4.5 Effects on In-and-Outs' Income

In this section I investigate whether in-and-outs appear to pay permanent costs for taking a break out of the labor force. Prior studies of unemployment have shown that workers subject to mass layoffs can suffer lower income for many years after the initial separation, but in-and-outs may not face the same consequences. If in-and-outs have job opportunities available, but choose to delay these opportunities while taking a break in between jobs, they may not experience any wage cut. However, they may not be able to keep up with the wage growth of their continuously employed peers, who will be accumulating human capital on the job while in-and-outs are out of the labor force.

I examine how in-and-outs' income changes relative to their peers before and after they leave the labor force. Using data from the SIPP, I examine changes in personal income of in-and-outs before and after a transition from employment to non-participation.

Figure 10: Personal Income of In-and-outs around Labor Force Transitions

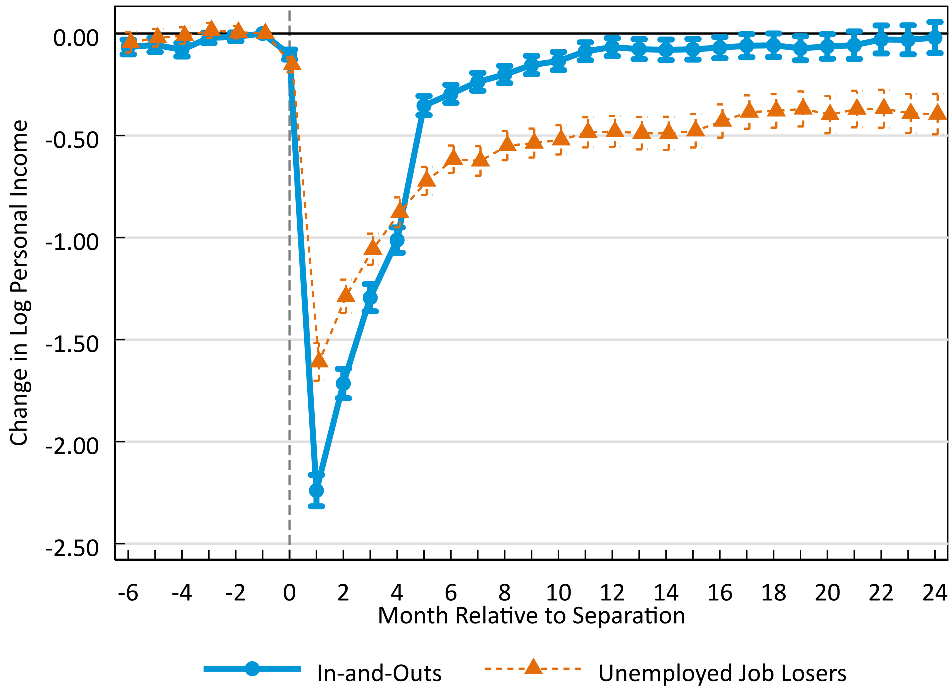

Source: SIPP 1984-2008, men ages 25-54. In-and-outs are employed in month 0, non-participants in month 1, and are employed again by at least month 12. Unemployed job losers are employed in month 0 and fired or laid off but looking for work in month 1. For each group, I plot the coefficients from estimating equation 4 over $k = -6, \ldots, 24$, where k refers to a monthly frequency. Since I control for year fixed effects and a cubic in age, the change in log personal income is relative to continuously employed individuals of the same age in the same year. Standard errors are clustered at the individual level. Personal income includes labor income, business income, transfers, and other income attributed to the respondent, but not income of other family members.

I use the same set of regressions as in equation 4 above, varying the horizon k from -6 to 24 in order to cover a two-and-a-half-year window around the labor force exit. As above, I control for a cubic in age as well as time-period fixed effects. I repeat this procedure using unemployed job losers in order to provide a baseline comparison.

Once back in the labor force, in-and-outs' earnings catch back up to that of their continuously employed peers. Figure 10 shows the estimated coefficients, which demonstrate that personal income relative to the continuously employed is the same 10 months after leaving the labor force as it was before the labor force exit. This is around the time when most in-and-outs have returned to the labor force. In comparison, the earnings of individuals who become unemployed fail to fully recover even 18 months later.

One potential concern with this analysis is that it may miss individuals who left the labor force intending to return quickly, but who subsequently received negative wage shocks and remained out of the labor force. In this way, the true cost of leaving the labor force is not identified by the above event study, since the sample is selected based on whether individuals returned to the labor force ex-post and this excludes individuals who receive negative wage shocks while out of the labor force. To address these concerns, I construct a sample of men who are predicted to be in-and-outs ex-ante, without using any ex-post information to select the sample. I employ a machine learning algorithm to predict which men will return to the labor force quickly after a separation based only on their pre-separation characteristics. Since this sample is defined using only ex-ante characteristics, the concern described above does not apply to this group. When I re-conduct the event study analysis on this sample, I find that this group of predicted in-and-outs experiences similar income dynamics as the sample of realized in-and-outs, indicating that the recovery of in-and-outs' earnings is not driven by sample selection problems. Full details of this analysis and the results are contained in Appendix E.

5 Alternative Explanations

The decline of prime age male labor force participation is commonly attributed to diminished labor demand or lower work incentives from taxes and transfers, but this section shows that neither of these forces can explain the rise of in-and-outs. I start by examining the wages available to in-and-outs, since lower wages could induce men to work less, but the rise of in-and-outs appears to have occurred at constant wages. Based on conventional estimates of labor supply elasticities, these changes in wages cannot explain a decrease in participation of the magnitude observed. The rise of in-and-outs has occurred across all industries and occupations, suggesting that this phenomenon is not related to particular types of jobs. Additionally, while wage rigidities can lead labor demand shocks to result in a temporary rise of in-and-outs, these effects disappear in the long-run. Lastly, I show that tax rates faced by in-and-outs have decreased slightly on average and that few in-and-outs receive government transfers, implying that neither explicit nor implicit tax rate increases could be behind the rise of in-and-outs.

5.1 Wages

Declining wages could be an important driver of diminished participation among men. Over the last few decades, several forces have emerged that could reduce the available wages for men, including skill-biased technological change, declining manufacturing, and decreased union coverage. Falling wages could induce men to spend less time working and take more time out of the labor force. Additionally, if wages fall to just above mens' reservation wage, small temporary shocks to productivity could lead men to leave the labor force for a short time before returning to work.

Figure 11: Wages and Participation Across Skill Levels

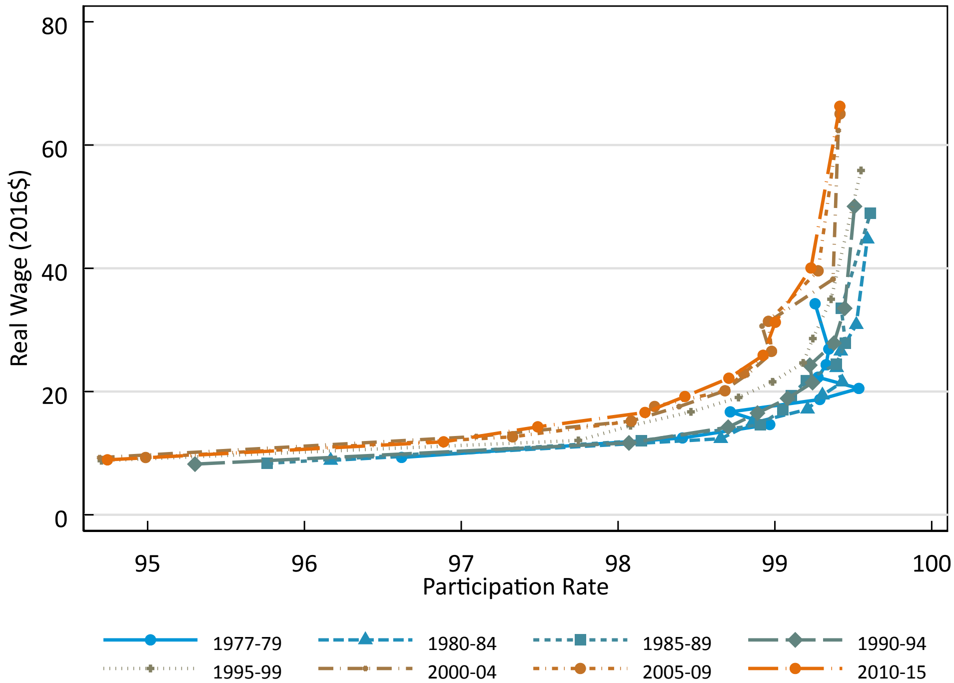

Source: CPS men ages 25-54 excluding dropouts, 1977-2015. Individuals are assigned to a wage decile based on the annual wage distribution in their year. For each five-year period, the average participation rate and average real wage are computed for each decile and plotted, pooling individuals across years within the five-year period. Nominal hourly wages are deflated by PCE price index. All statistics are computed using survey weights.

Fortunately, I can directly observe the wages of in-and-outs and examine how they have changed over time. I compute real hourly wages from the CPS Outgoing Rotation Groups for all prime age men, adjusting for top coding.[33] Wages for some in-and-outs are missing if they happened to be out of the labor force during the month of the ORG interview, but since CPS cohorts are randomly selected, this meets the missing-at-random criterion so dropping these observations will not confound the estimates (Rubin, 1976).[34] The full details describing how I construct wage data are contained in Appendix B.

Across the skill distribution, participation has fallen holding market opportunities constant. In Figure 11, I divide men into wage deciles based on their rank in the annual wage distribution and plot average participation (excluding dropouts) against average wages within each decile, repeating this separately for several eras. In-and-outs have

[33]Nominal wages are deflated by the PCE price index. I impute top-coded wages with the average above the top-code using a log-normal approximation following the method of Schmitt (2003).

[34]Alternatively, these missing wages can be imputed from the wages of other in-and-outs with similar labor force attachment, but the results are nearly identical under several different imputation strategies.

become more common at every wage level, resulting in this curve shifting to the left over time. This shift is more pronounced at the bottom of the skill distribution, where real wages have slightly declined, but higher skill levels have seen rising in-and-outs as well, even as wages have increased. Although there are some groups that have experienced declining wages, the leftward shift of this curve at every wage level makes it unlikely that changing wages are the main factor responsible for the rise of in-and-outs.

It is worth asking how much of the rise of in-and-outs can be attributed to the observed changes in wages of prime age men. The response of labor supply to changes in wages over the long-run is summarized by the uncompensated, or Marshallian, labor supply elasticity. Given the observed change in wages, the portion of the decline in participation attributable to labor demand can be expressed as

$$\Delta \log \widehat{\text{LFPR}} = \underbrace{\varepsilon_{l,w}}_{\substack{\text{Uncompensated} \\ \text{Labor Supply Elasticity}}} \cdot \underbrace{\Delta \log \frac{w}{p}}_{\substack{\text{Change in} \\ \text{Real Wages}}}$$

To measure the uncompensated elasticity, I average estimates of this elasticity from a long literature on male labor supply.[35] The average of these estimates is 0.04, although the estimates range from a minimum of -0.02 to 0.14 at maximum.[36] The elasticities used here measure the response of total labor supply to wages, including both in-and-outs and dropouts. Therefore, these can be thought of as an upper bound for the relevant labor supply elasticity, which would measure the response along the in-and-out margin alone.

Changes in wages explain very little of the rise of in-and-outs. Table 7a shows the predicted change in participation implied by the total change in wages and several possible values of the uncompensated labor supply elasticity. Given the 16.6% increase in real wages between 1977 and 2015, the central elasticity estimate from the literature of 0.04 predicts that changes in wages alone would lead to a 0.7% increase in labor supply, or a 0.6 p.p. increase in the participation rate.[37] For higher values of the uncompensated

[35] Specifically, I take the elasticity estimates from Hall (1970); Hausman (1981); Pencavel (1986); MaCurdy, Green and Paarsch (1990); Triest (1990); Juhn et al. (1991); Ziliak and Kniesner (1999); Juhn et al. (2002); Pencavel (2002); Eissa and Hoynes (2004); and Moffitt (2012).

[36] This range is lower on average than the range of labor supply elasticities highlighted in the meta-analyses of Chetty, Guren, Manoli and Weber (2013) and Chetty (2012), primarily due to the fact that I am focusing on prime age men in the United States, which excludes some of the higher elasticity estimates included in those meta-analyses.

[37] The average wage for employed prime age men rose 16.6% between 1977 and 2015, but this may overstate the rise in potential wages if lower wage individuals are increasingly not employed. When wages of the non-employed are imputed using the procedures outlined in Juhn (1992) and Blau and Kahn (2007), the increase in overall average real wages falls to 10.6% and 8.7% respectively, suggesting that some selection is occurring. However, when imputing wages of the non-employed using the selection correction of Heckman (1979), the increase in average real wages is 39.5% instead. Regardless, under all

Table 7: Contributions of Changing Wages to Participation Decline, 1977-2015

(a) Constant Elasticity

Elasticity Value	ΔLFP
0.04	0.63
-0.05	-0.78
0.0	0.00
0.1	1.56
0.2	3.13

(b) Heterogeneous Elasticities

Elasticities Source	ΔLFP
Juhn, Murphy, Topel (1991)	1.06
Pencavel (2002)	-3.05

Source: CPS men ages 25-54, 1977-2015. Predicted percent change in LFP is computed by multiplying the observed change in average wages by the given elasticity, which is then converted into the predicted percentage point change in LFP, which is reported above. For the heterogeneous elasticity calculations, this procedure is conducted separately for each decile of the wage distribution and the resulting predictions are aggregated to form an estimate of the overall predicted change. Wage changes are computed using survey weights.

elasticity an even larger increase in participation would be predicted. Only under a negative elasticity, indicating that income effects are substantially larger than substitution effects, would the observed changes in wages predict a decline in labor supply.

Table 7b relaxes the assumption of an identical labor supply elasticity for every individual, allowing elasticities to vary across the skill distribution. I take estimates of the elasticities at different skill levels from Juhn et al. (1991) and Pencavel (2002). I assign individuals to skill deciles based on their rank in the annual wage distribution and compute the change in average wages within each skill decile. Multiplying this change in wages by the elasticity for the decile and aggregating yields an estimate of the overall predicted change in participation. Using the elasticity estimates of Juhn et al. (1991), I estimate that the participation rate would have risen by 1.06 p.p. between 1977 and 2015, due to a small decline in participation for the lower half of the skill distribution being more than offset by an increase in participation among the top half.[38] In contrast, the elasticities estimated by Pencavel (2002) imply a net reduction in participation of 3.05 p.p. over this time period, mainly driven by large income effects among the top two deciles. Neither of these approaches matches the observed pattern of rising in-and-outs across deciles, with in-and-outs rising among every wage decile and larger increases at the bottom end of the distribution.

of these approaches male labor supply would be predicted to grow between 1977 and 2015 using the central elasticity estimate from the literature. More details on the imputation procedures can be found in Appendix B.

[38] This calculation uses average wages among employed prime age men without imputing wages to the non-employed. If wages are imputed using the methods of Juhn (1992) and Blau and Kahn (2007), the predicted changes are reduced to 0.45 p.p. and -0.45 p.p. respectively as both of these methods estimate larger declines in wages among the bottom half of the skill distribution relative to the no-imputation baseline. If wages are imputed with the selection correction of Heckman (1979), the predicted change rises to 4.7 p.p. instead as wages are estimated to have risen among the bottom half of the distribution. More details on the imputation procedures can be found in Appendix B.

Table 8: Rise of In-and-Outs Across Occupations

Occupation Group	ΔIn-and-Outs, 1980-2015 (p.p.)
High Abstract Tasks	2.2
High Routine Tasks	3.7
High Manual Tasks	3.7

Source: IPUMS CPS, 1980-2015. In-and-outs are defined as men ages 25-54 with between 1 and 7 months in the labor force out of 8 total. Individuals are categorized to occupations based on the occupation of their job when they are employed. Individuals who are not employed in any of their 8 CPS responses are excluded. For task content groups, I take occupations in the highest third of task content for each task type, where task content is measured as in Autor, Levy, & Murnane (2003). For each occupational category, I report the change in the share of individuals within that category who are in-and-outs between 1980 and 2015. All statistics are computed using survey weights.

5.2 Industries and Occupations

The increase of men cycling in and out of the labor force has happened in every industry, despite large differences in how industries have evolved over this time period. I take in-and-outs in the CPS and categorize them based on the industry (or industries) they work in while they are employed. Appendix Table A.4 shows the change in the share of employment made up by in-and-outs within each industry from 1980 to 2015. All industries have seen at least some increase in the share of in-and-outs and across most industries the shares of in-and-outs have risen between 3 and 7 p.p. over this time period. Industries with very different types of jobs have nonetheless seen similar increases in in-and-outs, as well as industries on different trajectories. The common growth of in-and-outs across a wide array of industries with different work patterns and different prospects for the future suggest that the rise of in-and-outs may not be strongly related to characteristics of jobs.

In-and-outs have grown as a share of many different types of occupations. Table 8 shows the change in the in-and-out share of employment among occupations involving different types of tasks. Among occupations with high abstract task content, the in-and-out share has risen by 2.2 p.p., only slightly less than the 3.7 p.p. seen in high routine and manual task occupations. Appendix Table A.5 further shows that in-and-outs have risen across six major groups of occupations, providing more evidence against the notion that the rise of in-and-outs is related to particular types of jobs. Since the rise of in-and-outs has occurred similarly across many different types of jobs, this suggests that changes in labor demand are unlikely to explain most of the rise of in-and-outs.

5.3 Shocks to Employment Opportunities

Changing labor demand could alternatively result in lower opportunities for employment without changing wages. For example, if wages are rigid, a fall in productivity may induce employers to ration jobs, which would result in fewer men being employed without a drop in wages. Blanchard and Katz (1992) provide evidence that regional labor demand shocks produce responses in employment and wages consistent with this pattern, although changes in participation in response to these shocks are found to be largely temporary.

I test the contribution of this force using regional shocks to labor demand. Shocks may have long-term effects on participation, which are captured by using an event study design to examine responses to shocks over a long horizon. I regress the change in the in-and-out share over different horizons on shocks to employment and controls:

$$\Delta\text{In-and-Outs}_{s,t,t-k} = \beta^{(k)}\Delta E_{s,t,t-1} + X_{s,t} + \epsilon_{s,t} \tag{5}$$

I repeat this regression for values of k between -4 and 5 to examine the response of in-and-outs to employment shocks in the five years before and after the shock. State-level in-and-out shares are computed from the CPS, and state-level employment growth is taken from the Quarterly Census of Employment and Wages (QCEW).[39] In the baseline, I control for time fixed effects alone, but since I conduct the regression in differences, this implicitly controls for state fixed effects as well. Standard errors are clustered at the state level.

To avoid potential endogeneity of employment changes, I instrument for the actual employment growth with predicted employment growth based on industry composition.[40] I construct the predicted change in state s and year t by combining national growth rates of industries i with lagged local industry shares:

$$\Delta\hat{E}_{s,t,t-1} = \sum_i (\log E_{-s,i,t} - \log E_{-s,i,t-1}) \cdot \frac{E_{s,i,t-3}}{E_{s,t-3}} \tag{6}$$

[39] The QCEW employment count includes jobs worked by non-prime-age-men. One concern could be that prime age men are not affected by general shocks to labor demand, making this a poor measure of labor demand for prime age men. However, employment of prime age men as measured by the CPS responds one-for-one to shocks to QCEW employment, allaying this concern. Another data source, the Quarterly Workforce Indicators (QWI), allows one to break out prime age male employment by state and industry separately, but unfortunately many of these data points are censored for confidentiality purposes and so this data source produces very noisy estimates. For this reason, I present estimates based on the QCEW data series.

[40] I use the approach of Bartik (1991) and Blanchard and Katz (1992) to compute predicted employment growth, but this is similar to using local industry shares directly as instruments (Goldsmith-Pinkham, Sorkin and Swift, 2017).

Figure 12: Response of In-and-Out Share to Labor Demand Shock

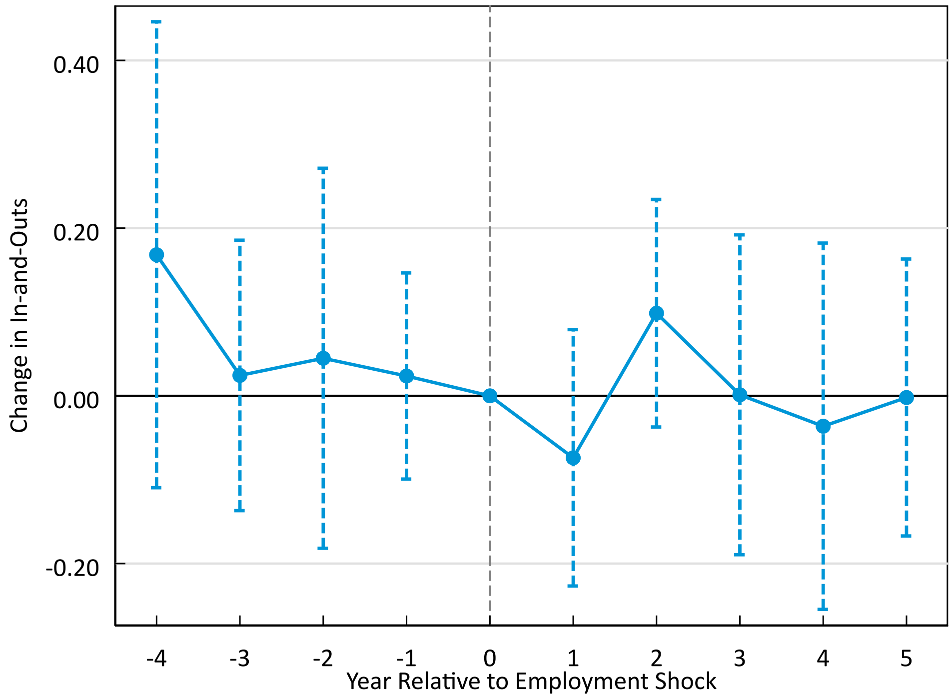

Source: State-level in-and-out share computed from IPUMS CPS using survey weights. Employment and Bartik instrument computed from QCEW. Using equation 5, the change in the state-level in-and-out share over a horizon k is regressed on the growth rate of employment instrumented with a Bartik measure of predicted employment growth. By plotting the coefficients $\beta^{(k)}$, this traces out the impact of an exogenous 1% increase in state-level employment on the in-and-out share. The Bartik measure is computed as described in equation 6 using a leave-one-out procedure. The first stage F statistic is above 10 for every horizon.

By using national growth rates excluding state s to predict the change in employment in state s, this avoids mechanical correlation between predicted and actual employment growth. I take data on employment by state and industry from the QCEW.

Employment shocks appear to have little effect on temporary non-participation in the form of in-and-outs. Figure 12 plots the $\beta^{(k)}$ coefficients from the regression above to show the dynamic response of in-and-outs around a positive 1% shock to employment growth. Immediately after the shock, the in-and-out share appears to drops slightly by 0.08 p.p., although a response of zero cannot be ruled out. The immediate response does not seem to be very persistent, though, as the in-and-out share appears unchanged five years after the shock. The confidence intervals rule out an increase or decrease of more than 0.2 p.p. in the five years after a 1% employment shock. This indicates that shocks to labor demand holding wages constant explain little to none of the rise of in-and-outs.

5.4 Taxes & Transfer Programs

Changes in taxes and transfer programs can have significant effects on labor force participation. Increases in marginal tax rates reduce the returns to work and could result in men taking more time out of the labor force. This applies both to explicit marginal tax rates from income taxes as well as the implicit marginal tax rates from means-tested transfers phasing out. Some programs, such as Social Security Disability Insurance (SSDI), phase out completely when a recipient rejoins the labor force. Prior work has documented increases in the tax rates faced by US households, which could result in lower labor force participation (Mulligan, 2002).

Between the late 1980s and 2015, income tax rates have actually declined somewhat for prime age men. Using a sample from the Basic CPS matched to income information from the March CPS Supplement and taxes computed from NBER's TAXSIM calculator, I estimate the distribution of income tax rates faced by prime age men, including federal, state, and payroll taxes.[41] Appendix Figure A.2 shows that average tax rates for in-and-outs have fallen by nearly half over this period, from slightly more than 15% in the late 1980s to around 10% in 2015, while marginal tax rates are roughly constant at around 25% or perhaps slightly declining. Tax rates are countercyclical due to a progressive tax code; as incomes fall during recessions, tax rates fall accordingly. Beyond the cyclical variation, both average and marginal tax rates exhibit steady downward trends. On the whole, income tax rates for in-and-outs fell substantially over this time period, implying that the rise of in-and-outs cannot be attributed to rising income tax rates.

In-and-outs have been largely unaffected by changes to Social Security Disability Insurance (SSDI) program, one of the largest cash-based transfer programs. Partially due to changes in SSDI eligibility policies, the number of SSDI beneficiaries has risen dramatically in the last several decades, nearly tripling between 1980 and 2015. This program imposes very large implicit marginal tax rates for recipients, since eligibility for the program expires when a worker becomes re-employed, causing some to attribute declining prime age male labor force participation to changes in this program (Autor and Duggan, 2003). While this may have contributed to the increase of dropouts, it cannot explain the rise of in-and-outs. Fewer than 7% of in-and-outs live in a household receiving any SSDI benefits in the March CPS. Additionally, most in-and-outs are not out of the labor force long enough to apply for benefits, have their case examined, and start receiving benefits before rejoining the labor force.

Furthermore, most in-and-outs receive no cash-based transfers of any kind. Using data from the Survey of Income and Program Participation (SIPP), I investigate the share of prime age men living in families that receive any means-tested cash transfers, including General Assistance, Supplemental Security Income, Temporary Assistance for

[41]More details on the construction of these tax rates are available in Appendix B.

Figure 13: Transfer Receipt by Labor Force Participation Category

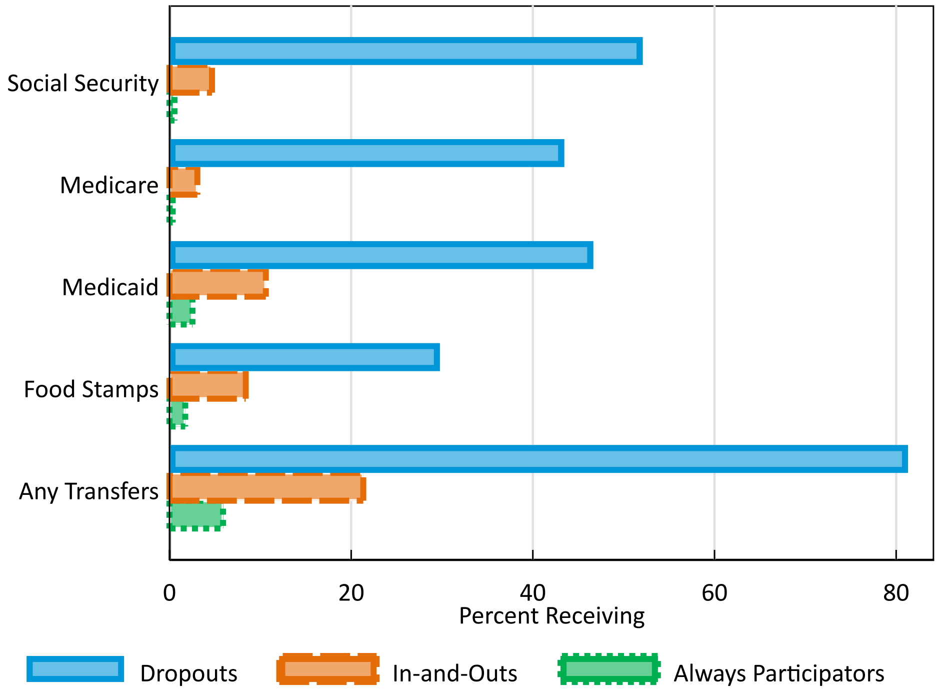

Source: SIPP men ages 25-54. Transfer receipt is computed on a monthly basis. Individuals are denoted as receiving a particular form of transfers in a month if anyone in their family reports receiving this form of transfers. The category of "Any Transfers" includes in-kind transfers and means-tested cash-based transfers. All statistics are computed using survey weights.

Needy Families, and other programs. Figure 13 shows that less than 10% of in-and-outs' families received any income from means-tested cash transfers, compared to nearly 40% for dropouts. Prior work has pointed out that the SIPP may suffer from underreporting of participation in some of these programs (Meyer, Mok and Sullivan, 2015), but even if the true rate were double the measured rate, more than 80% of all in-and-outs would be receiving no transfer income.

In-and-outs are also unlikely to be receiving in-kind transfers, such as food stamps or Medicaid. While nearly half of dropouts receive Medicaid benefits, as shown in Figure 13, only about 10% of in-and-outs receive these benefits and even fewer in-and-outs receive food stamps or Medicare benefits. Of those receiving some form of in-kind transfers, many in-and-outs receive more than one form of in-kind transfers, such that only about 17% of in-and-outs receive any form of in-kind transfers. Combining this with cash transfers as mentioned above shows that only about 20% of in-and-outs receive any form of transfers. With such a low rate of transfer receipt, it is unlikely that the implicit taxes created by transfer programs can explain much of the rise of in-and-outs.

6 Model

This section combines the channels identified in previous sections within a unified framework. Reduced-form evidence suggests that changes both within and across households account for the rise of in-and-outs, so I begin by laying out a dynamic model of labor supply and household formation to capture these two channels. This model allows for changes in participation within households coming from both supply and demand forces, such as those analyzed previously, as well as changes induced by reallocation across household types. The model is fit to data on participation and household types over the last several decades to estimate unobserved parameters. I estimate the model from 1978 to 2015 and measure the contribution of changes in each parameter to the rise of in-and-outs. The model estimates imply that labor supply factors are responsible for nearly the entire rise of in-and-outs, with little contribution from changes in labor demand.

6.1 Setup

This section outlines the setup of a dynamic model of labor supply and household formation and its solution. As in Chiappori (1992), I use a collective model of households in which individuals' outside options influence the allocation within households. Individuals face the choice of whether or not to move in with their parents, similar to the model introduced by Kaplan (2012), as well as whether to get married, as in Knowles (2012).[42]

The model consists of a continuum of men and women facing the same problem. Each period, individuals choose an amount of consumption c, labor supply $n \in [0,1]$, and a household type $h \in \{\text{Alone, Living with Parents, Married}\}$. I assume that individuals have common preferences over consumption and labor supply given by the utility function $u(c,n)$.

Labor supply is assumed to represent changes in labor force participation along the in-and-out margin only. Instead of modeling labor supply as binary within many small periods, I allow n to be continuous between zero and one, representing the fraction of time across some longer period in which an individual is in the labor force. I treat the periods in the model as years, meaning that n is equal to the fraction of the year that an

[42] Relative to this literature, the model used in this paper stands out for allowing individuals to choose over more than two household types as well as allowing for a rich set of shocks to preferences over household type. The model of Knowles (2012) involves individuals choosing between marriage and living alone with a single shock to preferences, while the model of Kaplan (2012) involves individuals and parents playing a repeated game in which the child can choose to move home and the parents choose an amount of resources to share with their child.

individual participates and $1 - n$ is the fraction of the year that the individual is out of the labor force as an in-and-out. Accordingly, changes in n are assumed to come solely from the in-and-out margin.[43]

Each period individuals choose their household type out of the available options. Since consumption and labor supply may vary across household types, individuals' potential utility will vary across their possible choices, although this variation is the same for all individuals. To introduce a role for idiosyncratic differences in utility across potential household types, I allow for an additively-separable, individual-specific utility component ω_i^h, which represents the benefit or cost to individual i of living in household type h. Combining these components, individual preferences are given by:

$$U_i^h = u(c, n) + \omega_i^h$$

Individuals receive a new i.i.d. draw of $\Omega_i \equiv \begin{pmatrix} \omega_i^A \\ \omega_i^{LWP} \\ \omega_i^M \end{pmatrix}$ each period.[44]

Income comes from two sources. First, individuals receive labor income for working at a gender-specific wage rate w_g. This assumption focuses the model on capturing the effects of differences in average wages between genders on labor supply, while omitting a role for heterogeneity in wages within each gender. This is a somewhat restrictive assumption, since it ignores changes in moments of the distribution of wages beyond the mean. Nevertheless, the analysis of changing wages by decile in section 5.1 reveals that capturing changes across the wage distribution does not help explain the rise of in-and-outs, so to a first-order this dimension is not important for the model to include.[45] Additionally, accounting for wage heterogeneity within genders would require that the household formation choice to be conditional on the wages of potential partners and incorporating strategic matching of this type is outside the scope of the present paper.

The second form of income individuals receive is unearned income Y_h, representing asset income or other forms of non-labor income, which I allow to vary across household

[43]When taking the model to the data, I exclude dropouts in calculating the empirical moments to be consistent with this assumption.

[44]The assumption that the random utility component is i.i.d. is necessary for tractability, as otherwise the household formation choice by individuals would need to be conditional on the past distribution of random utility components across the population of potential matches.

[45]With the assumption of a common wage rate for each gender, the model also cannot speak to whether individuals change household type in response to negative labor market shocks. This channel is important in the model of Kaplan (2012), as some individuals may use the option of moving back in with their parents as a form of insurance against negative shocks. However, while the wages of prime age men differ across household types, these differences are not large enough to explain the observed differences in labor force participation using conventional labor supply estimates from the literature, so this channel is unlikely to explain much of the rise of in-and-outs.

types. For individuals living with their parents, this unearned income may represent a transfer of income or a public good like housing provided within the household.[46]

Consumption and Labor Supply For individuals living alone or with their parents, the problem takes a straightforward form. Individual i maximizes utility subject to a budget constraint:

$$\begin{aligned} \max_{c,n} U_i^h &= u(c,n) + \omega_i^h \\ s.t.\, c &= w_{g(i)} n + Y_h \end{aligned} \tag{7}$$

Married households follow a collective model, as in Chiappori (1992). I assume that the allocation within a married household maximizes a weighted sum of individual utilities given a Pareto weight for each individual. Formally, for a household with individuals i and j, consumption and labor supply are determined by:

$$\max_{c_i, n_i c_j, n_j} \mu_{ij} U_i^M + (1 - \mu_{ij}) U_j^M \tag{8}$$

$$s.t.\, c_i + c_j = w_{g(i)} n_i + w_{g(j)} n_j + Y_M$$

where $U_i^M = u(c_i, n_i) + \omega_i^M$ and U_j^M is defined analogously for j.

The additive structure of utility within a married household disallows any complementarities in leisure between married individuals. To test the plausibility of this assumption, I repeat the event study design employed in Sections 4.4 and 4.5 with the labor supply of in-and-outs' partners as the outcome of interest. I find only a slight increase in the participation of in-and-outs' partners, suggesting that leisure complementarities are close to negligible, at least in the short-run.

Household Formation Each period, individuals must choose which type of household to live in. Single individuals are matched randomly with a partner of the opposite gender and must decide whether to marry or not. If they choose not to marry, then they choose whether to live alone or with their parents. Married individuals must choose whether to remain married or get a divorce, and if the latter they then each choose whether to live alone or with their parents. Divorce incurs a fixed cost k for each individual.

This setup gives rise to threshold rules governing the choice of housing type. For example, an individual will choose to live with his or her parents if the random utility

[46] I do not formally model the parents' decision, but the quantity of this transfer could be taken as the outcome of a dynamic game between parents and their children where parents have some altruistic preferences, as in Kaplan (2012).

component of preferences ω_i^{LWP} is above a threshold $\bar{\omega}^{LWP}(\mu,\Omega;h)$, where μ is the potential Pareto weight he or she would face if married.[47] The choice of household type can be summarized with a household decision rule for each gender $\Lambda_g^h(\mu,\Omega)$.

Value Functions Given a value for the Pareto weight μ, the value function for an individual of gender g living in household type h with random utility draw Ω can be written as:

$$\begin{aligned} V_g^h(\mu,\Omega) = U_g^h(\mu,\Omega) + & \\ \beta\left[\sum_{h'} \Pr(h' = \Lambda_g^h(\mu,\Omega))\mathbb{E}_{\mu',\Omega'|h'}\left(V_g^{h'}(\mu',\Omega') - k\cdot\mathbf{1}(\text{Divorce})\right)\right] & \qquad (9) \end{aligned}$$

Note that the dependence of V on μ is only relevant for married households.

Bargaining Solution I now turn to the process by which the Pareto weights μ are determined. Since individuals can choose household type each period, I allow the Pareto weight to adjust in response to changes in the outside option of each individual, consistent with the evidence on labor supply responses presented by Chiappori, Fortin and Lacroix (2002). Let the gains from marriage be given by:

$$S_i(\mu,\Omega) = V_{g(i)}^M(\mu,\Omega) - \max\{V_{g(i)}^S(\Omega), V_{g(i)}^{LWP}(\Omega)\} - k$$

Given these gains for each individual, I assume that μ_{ij} is determined through Nash bargaining with equal bargaining weights. This is equivalent to:

$$\mu(\Omega_i,\Omega_j) = \arg\max_{\mu} S_i(\mu,\Omega_i)\cdot S_j(1-\mu,\Omega_j) \qquad (10)$$

As a result of this setup, marriage and divorce are always mutually agreed on. If there exists a Pareto weight μ such that both individuals prefer the allocation with this weight to their outside options, then the individuals will agree to be married with this bargaining weight rather than divorce. Equivalently, married individuals will only divorce if there exists no weight μ and associated allocation that makes both individuals prefer marriage to their outside option.

Definition 1. *A* stationary recursive equilibrium *of this model consists of value functions for each household type $V_g^h(\mu,\Omega)$, a bargaining solution $\mu(\Omega_i,\Omega_j)$, household decision rules $\Lambda_g^h(\mu,\Omega)$, and consumption and labor supply policies $c^h(\mu)$ and $n^h(\mu)$ such that:*

[47] The dependence of this threshold on household type h is only due to the fixed cost of divorce. As a result of this, the threshold is the same for individuals living alone or with their parents, but is different for married individuals.

1. *The value functions $V_g^h(\mu, \Omega)$ solve equation 9;*
2. *The bargaining solution $\mu(\Omega_i, \Omega_j)$ is equal to the Nash bargaining solution of equation 10;*
3. *The household decision rule maximizes individuals' value, i.e.*

$$\Lambda_g^h(\mu, \Omega) = \arg\max_{h'} V_g^{h'}(\mu, \Omega)$$

4. *Consumption and labor supply policies of single individuals solve equation 7 and those of married individuals solve equation 8.*

6.2 Estimation

I now turn to estimating the parameters of this model over time. Within each time period, some of the parameters of this model can be estimated directly from the data while others need to be estimated indirectly by matching the model to empirical moments. I conduct indirect inference on the latter set of parameters using the Simulated Method of Moments (SMM) separately for each year as described below.

Functional Forms I make several assumptions regarding the functional forms of the model. First, I assume that the common utility function over consumption and labor supply takes on a standard balanced-growth form,

$$u(c, n) = \log(c) - \psi \frac{n^{1+\frac{1}{\theta}}}{1 + \frac{1}{\theta}}$$

where θ is the Frisch, or consumption-constant, elasticity of labor supply and ψ captures forces that shift the marginal disutility of work relative to the marginal utility of consumption. This choice is somewhat conservative since these preferences balance substitution and income effects and allow for a stable path of labor supply with economic growth, while some other possible preferences feature income effects that dominate substitution effects and imply that labor supply must decrease with economic growth Boppart and Krusell (2016).

I assume the random utility component Ω follows a multivariate normal distribution. Without loss of generality, I rewrite the random utility component Ω_i with three components in terms of the differences between shocks, yielding a transformed vector $\tilde{\Omega}_i$ with only two components:

$$\tilde{\Omega}_i \equiv \begin{pmatrix} \omega_i^{LWP} - \omega_i^S \\ \omega_i^M - \omega_i^S \end{pmatrix}$$

Table 9: Calibrated Parameters

Parameter	Value	
	1978-79	2014-15
w_m	1	1.15
w_f	0.574	0.945
Y_A	0.05	0.06
Y_M	0.03	0.01
β	0.985	

Source: IPUMS CPS, 1978-79 and 2014-15. Hourly real wages are normalized relative to male wages in 1978-79. Unearned income is total household income minus own labor income (and spousal labor income for married households). Unearned income in each period is normalized relative to annual income from a full-time full-year male worker in 1978-79. Empirical quantities are computed using survey weights.

Each component represents the net utility benefit or cost of a household type relative to living alone. Since the shocks are generated by a multivariate normal distribution, $\tilde{\Omega}_i$ also follows a multivariate normal distribution:

$$\tilde{\Omega}_i \sim N\left(\begin{pmatrix} \alpha^{LWP} \\ \alpha^{M} \end{pmatrix}, \begin{pmatrix} \left(\sigma^{LWP}\right)^2 & \rho\sigma^{LWP}\sigma^{M} \\ \rho\sigma^{LWP}\sigma^{M} & \left(\sigma^{M}\right)^2 \end{pmatrix}\right)$$

Calibrated Parameters The calibrated values of several parameters are reported in Table 9. I estimate the wages for each gender, w_g, directly from the data, normalizing all quantities relative to the wages of men in 1978-79, which I set to be equal to 1. Additionally, I estimate the median levels of unearned income Y_h for individuals living alone and married individuals directly from the data, normalizing these values relative to annual income of a full-time full-year male worker in 1978-79. This measure includes all income from transfers, assets, and businesses within the household.

Estimated Parameters The remaining parameters are all estimated via Simulated Method of Moments (SMM). These include the preference parameters ψ and θ, the mean and variance parameters of the random utility component, the divorce cost k, and the amount of unearned income for individuals living with parents Y_{LWP}.[48]

For each time period, the model is matched to three sets of moments. First, the model is required to match the share of men living in each household type. Second, the simulated labor supply within each household type is matched to its empirical counterpart. Lastly, I set the model to match the annual transition rates between household types, e.g. the probability of presently living in a married household conditional on

[48] While unearned income for other household types is observed, the amount of parents' income that is shared with their children must be estimated indirectly.

having lived with parents one year prior.[49] All moments are computed from the CPS excluding dropouts, to be consistent with the assumption that labor supply in the model reflects in-and-outs only. I compute these moments pooling all CPS respondents that began their interviews in years 1978-79 and again for 2014-2015.

6.3 Results

I start by examining the fit of the model to the empirical moments, finding that it matches the distribution of household types and the participation rates quite well, while somewhat overestimating the flows between household types. Next, I use the estimated values of the parameters to measure the contribution of different supply and demand forces to the rise of in-and-outs. I find that growing female wages are responsible for about half of the rise of in-and-outs, consistent with the reduced-form evidence presented above, while the remainder is largely due to a combination of other labor supply forces, including transfers from parents.

Model Fit Table 10 shows the estimated parameters for 1978-79 and 2014-15 along with the values of the moments used in SMM. For the usual reasons, there is not a direct mapping between each parameter and corresponding moments, but Table 10 matches up parameters and the moments used to identify them approximately.

The model is capable of matching the shares of each household type and the participation rates within each household type quite well. The model estimates large average utility benefits of marriage, while individuals get a substantial utility cost from living with their parents on average. As the observed share of married men has declined over time, the model reflects this with a slight decline in the random utility component for marriage.

The pattern of simulated participation rates across household types is very similar to the observed pattern. Within each time period, participation is highest for married men and only slightly lower for men living alone, while men living with their parents participate at substantially lower levels. The model accounts for this by estimating that men living with their parents receive substantial transfers, between about one third to one half of average annual wage income. In both time periods, the estimated Frisch elasticity is about 0.4, which is only slightly higher than the meta-analysis estimate of Chetty et al. (2013) and within the bounds estimated by Chetty (2012).

While most of the moments are matched quite well, the model somewhat overestimates the flows between household types, particularly for flows into marriage from the

[49] Annual transition rates are measured in the CPS by comparing longitudinal matches one year apart, e.g. the first and fifth interviews or the second and sixth interviews, etc.

Table 10: Estimated Parameters

Parameter	Value		Moment	Simulated		Actual	
	1978-79	2014-15		1978-79	2014-15	1978-79	2014-15
ψ	1.05	1.07	n_A	97.3	96.6	96.7	95
θ	0.4	0.42	n_M	98.4	97.1	98.6	97.1
Y_{LWP}	0.29	0.45	n_{LWP}	91.3	88	92.1	88.3
α^{LWP}	-0.55	-0.52	s^{LWP}	13	19	5	12
α^M	0.68	0.64	s^M	78	67	85	69
σ^{LWP}	0.24	0.29	A→M	78	67	15	6.6
σ^M	0.21	0.24	LWP→M	78	67	1.2	1.3
ρ	-0.0003	-0.0001	A→LWP	9	14	0.014	0.024
k	0.19	0.2	M→LWP	9	14	0	0
			M→A	13	19	0	0
			LWP→A	13	19	0.1	0.28
			N	10,000,000	10,000,000	354,888	301,432

Notes: Parameters are estimated separately for each time period using SMM. The model is simulated for 100,000 individuals for 100 time periods, after a 10 period burn-in. The moments are computed pooling across all individuals and time periods. Empirical moments are computed from the monthly IPUMS CPS matched longitudinally for each time period. Empirical shares of household types do not add to 100%, due to "Other" category, which is missing from the model, but this difference is negligible. Participation rates within each household type are computed excluding dropouts, so variation in participation comes from in-and-outs only. Hazard rates for transitions between household types are computed at an annual frequency using individuals' household type during the last four months of the CPS rotation compared to household type from the first four months of their rotation. Empirical moments are computed using survey weights.

other two household types. This is a consequence of the assumption that the random utility component is i.i.d. rather than being somewhat persistent. However, a persistent shock would introduce strategic matching elements into the household formation decision, which I have elected to avoid in the interest of parsimony. While the simulated flows consistently overestimate the empirical flows, the overall pattern is very similar, with the largest hazard rates for transitions into marriage and lower hazard rates for the other types of transitions. Since the focus of the model is on participation rates, which are fit quite well, the overestimation of transition rates should not be a central concern.

Contributions The estimated model can be used to quantify the contributions of different forces to the decline in participation due to in-and-outs. In particular, I focus on the contributions of several different labor supply factors, including wages of men's partners, unearned income, and preferences, as well as the contribution of labor demand through changing male wages. For each of these factors, the model is able to estimate the total contribution to the participation, capturing both direct effects on participation within households as well as indirect effects through changing the distribution of house-

Table 11: Estimated Contributions

Parameters	Contribution
w_m	0.41
w_f	-0.92
Y_A, Y_{LWP}, Y_M	-0.43
α_{LWP}, α_M, σ_{LWP}, σ_M, ρ	-0.36
ψ	-0.4
θ	-0.11
k	-0.054
Residual	-0.23
Total Decline, 1978-2015	-2.1

Notes: Decline refers to the change in participation due to in-and-outs between 1978 and 2015. Contributions are measured by allowing one set of parameters to change over time in line with the estimated values from each time period, holding all other parameters constant. Residual contribution measures the difference between the simulated 1978-2015 decline in the model and the actual decline measured in the longitudinal monthly CPS.

hold types.

I estimate the contributions of several sets of parameters by examining the simulated change in participation holding all other parameters constant. Specifically, I begin by dividing the total set of parameters into Q non-overlapping sets $S^1, \ldots, S^Q$. I denote the set of parameters S^q evaluated at their estimated period t values as S^q_t. The participation rate in period t simulated by the model can be written as a function of the parameters in that period, i.e. $n(S^1_t, \ldots, S^Q_t)$. Using this setup, the total change in the simulated participation rate between period 0 and t can be decomposed into the contributions of changes in each parameter as follows:

$$\begin{aligned}\Delta n(S^1_t, \ldots, S^Q_t) &= n(S^1_t, \ldots, S^Q_t) - n(S^1_0, \ldots, S^Q_0) \\ &= \sum_q \underbrace{n(\ldots, S^{q-1}_t, S^q_t, S^{q+1}_0, \ldots) - n(\ldots, S^{q-1}_t, S^q_0, S^{q+1}_0, \ldots)}_{\text{Contribution of changes in } S^q}\end{aligned}$$

In this way, I examine the contributions of each set by feeding their changing values into the model, one set at a time, holding all other parameters constant.[50] I add to this a residual component equal to the difference between the simulated decline in participation and the actual decline, since the model is not able to match the participation rate in each period exactly. This yields a complete decomposition of the decline in participation due to in-and-outs into the contributions of different components, including supply and demand forces.

Table 11 reports the contribution of each set of factors to the decline in participa-

[50]The estimated contributions are not sensitive to the order in which I examine them, varying only slightly across different orderings.

tion due to in-and-outs. Consistent with the reduced-form evidence presented earlier, increased male wages are estimated to have offset the decline somewhat (contributing 0.4 p.p.), while increased female wages are responsible for slightly less than half of the decline (-0.9 p.p.). One fifth of the decline is attributed to changing unearned income (-0.4 p.p.) and an additional fifth is attributed to changes in preferences for household type (-0.4 p.p.). Most of the remainder is attributed to a shift in the disutility of labor represented by ψ (-0.4 p.p.).

The fact that the estimated contributions of changing wages match the reduced-form evidence indicates that the indirect effects of wages on household structure are minor. This is consistent with findings by Shenhav (2016) and Autor, Dorn and Hanson (2017) that changes in relative wages can explain less than one third of the decline in marriage. Instead of relative wages, the model mostly attributes the observed changes in household structure to the combination of changes in unearned income and preferences for household type. These forces could reflect underlying explanations related to the growth of men living with their parents, such as the effects of rising housing costs, which indirectly lower labor supply through the wealth effect from intergenerational transfers.

Additionally, the shift in the disutility of labor reflected by ψ suggests a small role for a labor supply shift that affects all households equally. This could include a change in the value of leisure, such as the increasing quality of video game technology proposed by Aguiar et al. (2017), or alternatively an increase in the disutility of work. Nevertheless, the model presented here suggests that such an explanation can only explain about one fifth of the total rise of in-and-outs.

7 Conclusion

Many different margins can contribute to a decline in labor supply. This paper has identified an understudied margin responsible for one-third of the participation decline among prime age men since the mid-1970s. The rise of in-and-outs forms a different picture of declining labor supply than the common view of men permanently withdrawn from the labor force. In-and-outs are highly attached to the labor force, work typical jobs, and are only notable in that they take brief breaks out of the labor force. I have presented evidence that the increasing prevalence of these breaks reflects a change in the desired amount of lifetime labor supply, due to several different labor supply factors, with little role for explanations related to labor demand.

Since in-and-outs' breaks are typically short and infrequent, as well as incurring no permanent costs, the rise of in-and-outs likely has had a minimal effect on the inequality of lifetime incomes among men. In contrast, the growth of permanent dropouts over the last several decades, responsible for the other two-thirds of the decline in participation since 1977, likely raised lifetime income inequality substantially. Dropouts not only forgo years of income while out of the labor force but may also experience lower wages if they return to the labor force, compounding the increase in inequality. The differences in consequences between changes along these two margins illustrates the importance of separating in-and-outs from dropouts when studying changes in labor supply.

To the extent that the rise of in-and-outs over the last several decades continues through the next several decades, this phenomenon could have important consequences for many aspects of the economy. For example, reduced labor force growth can slow the overall rate of economic growth, as occurred over the decade following the Great Recession (Fernald, Hall, Stock and Watson, 2017). Additionally, declining labor force participation may make jobless recoveries more common, as employment takes longer to reach its previous peak after a recession. The rise of in-and-outs may have social consequences too, particularly for families, as men more frequently take short breaks out of the labor force. This phenomenon merits continued study to understand the full consequences of rising in-and-outs.

References

Abowd, John M. and Arnold Zellner (1985) "Estimating gross labor-force flows," *Journal of Business & Economic Statistics*, Vol. 3, pp. 254–283.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst (2017) "Leisure luxuries and the labor supply of young men," NBER Working Paper.

Aguiar, Mark and Erik Hurst (2005) "Consumption versus Expenditure," *Journal of Political Economy*, Vol. 113, pp. 919–948.

Akerlof, George A. and Janet L. Yellen (1985) "Unemployment through the Filter of Memory," *The Quarterly Journal of Economics*, Vol. 100, pp. 747–773.

Autor, David H., David Dorn, and Gordon H. Hanson (2013) "The China syndrome: Local labor market effects of import competition in the United States," *The American Economic Review*, Vol. 103, pp. 2121–2168.

——— (2017) "When work disappears: manufacturing decline and the falling marriage-market value of men," NBER Working Paper.

Autor, David H. and Mark G. Duggan (2003) "The rise in the disability rolls and the decline in unemployment," *The Quarterly Journal of Economics*, Vol. 118, pp. 157–206.

Autor, David H. and Melanie Wasserman (2013) "Wayward sons: The emerging gender gap in labor markets and education," *Third Way Report*.

Barnichon, Regis and Andrew Figura (2016) "Declining desire to work and downward trends in unemployment and participation," *NBER Macroeconomics Annual*, Vol. 30, pp. 449–494.

Bartik, Timothy J. (1991) *Who Benefits from State and Local Economic Development Policies?*: WE Upjohn Institute for Employment Research.

Becker, Gary S. (1965) "A Theory of the Allocation of Time," *Economic Journal*, pp. 493–517.

Blanchard, Olivier Jean and Lawrence F. Katz (1992) "Regional evolutions," *Brookings Papers on Economic Activity*, Vol. 1992, pp. 1–75.

Blau, Francine D. and Lawrence M. Kahn (2007) "Changes in the labor supply behavior of married women: 1980–2000," *Journal of Labor Economics*, Vol. 25, pp. 393–438.

Boppart, Timo and Per Krusell (2016) "Labor supply in the past, present, and future: a balanced-growth perspective," NBER Working Paper.

Browning, Martin and Thomas F. Crossley (2001) "Unemployment insurance benefit levels and consumption changes," *Journal of Public Economics*, Vol. 80, pp. 1–23.

Chetty, Raj (2012) "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply," *Econometrica*, Vol. 80, pp. 969–1018.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber (2013) "Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities," *NBER Macroeconomics Annual*, Vol. 27, pp. 1–56.

Chetty, Raj and Adam Szeidl (2007) "Consumption commitments and risk preferences," *The Quarterly Journal of Economics*, Vol. 122, pp. 831–877.

Chiappori, Pierre-Andre (1992) "Collective labor supply and welfare," *Journal of Political Economy*, Vol. 100, pp. 437–467.

Chiappori, Pierre-Andre, Bernard Fortin, and Guy Lacroix (2002) "Marriage market, divorce legislation, and household labor supply," *Journal of Political Economy*, Vol. 110, pp. 37–72.

Council of Economic Advisers (2016) "The Long-Term Decline in Prime-Age Male Labor Force Participation."

Davis, Steven J. and John Haltiwanger (2014) "Labor market fluidity and economic performance," NBER Working Paper.

Drew, Julia A. Rivera, Sarah Flood, and John Robert Warren (2014) "Making full use of the longitudinal design of the Current Population Survey: Methods for linking records across 16 months," *Journal of Economic and Social Measurement*, Vol. 39, p. 121.

Eberstadt, Nicholas (2016) *Men Without Work: America's Invisible Crisis*: Templeton Foundation Press.

Eissa, Nada and Hilary Williamson Hoynes (2004) "Taxes and the labor market participation of married couples: the earned income tax credit," *Journal of Public Economics*, Vol. 88, pp. 1931–1958.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin (2015) "On the importance of the participation margin for labor market fluctuations," *Journal of Monetary Economics*, Vol. 72, pp. 64–82.

Elsby, Michael W. L. and Matthew D. Shapiro (2012) "Why does trend growth affect equilibrium employment? A new explanation of an old puzzle," *The American Economic Review*, Vol. 102, pp. 1378–1413.

Fernald, John G., Robert E. Hall, James H. Stock, and Mark W. Watson (2017) "The Disappointing Recovery of Output after 2009," *Brookings Papers on Economic Activity*.

Flinn, Christopher J and James J Heckman (1983) "Are unemployment and out of the labor force behaviorally distinct labor force states?" *Journal of Labor Economics*, Vol. 1, pp. 28–42.

Flood, Sarah, Miriam King, Steven Ruggles, and J. Robert Warren (2017) "Integrated Public Use Microdata Series, Curr," [dataset]. Minneapolis: University of Minnesota. https://doi.org/10.18128/D030.V5.0.

Flood, Sarah and José Pacas (2016) "Using the Annual Social and Economic Supplement with Current Population Survey Panels," Working Paper.

Ganong, Peter and Pascal Noel (2015) "How does unemployment affect consumer spending?", Harvard University Working Paper.

Goldin, Claudia (1989) "Life-cycle labor-force participation of married women: Historical evidence and implications," *Journal of Labor Economics*, Vol. 7, pp. 20–47.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2017) "Bartik instruments: What, when, why and how," Working Paper.

Greenwood, Jeremy, Nezih Guner, and Guillaume Vandenbroucke (2017) "Family economics writ large," NBER Working Paper.

Gruber, Jonathan (1997) "The consumption smoothing benefits of unemployment insurance," *The American Economic Review*, Vol. 87, pp. 192–205.

Hall, Robert E. (1970) "Wages, income and hours of work in the US labor force," MIT Working Paper.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda (2013) "Who creates jobs? Small versus large versus young," *Review of Economics and Statistics*, Vol. 95, pp. 347–361.

Hansen, Gary D. (1985) "Indivisible labor and the business cycle," *Journal of Monetary Economics*, Vol. 16, pp. 309–327.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*: Springer Science & Business Media.

Hausman, Jerry A. (1981) "Labor supply," in Henry J. Aaron and Joseph A. Pechman eds. *How Taxes Affect Economic Behavior*, Washington, D.C.: The Brookings Institution, pp. 27–72.

Heckman, James J. (1979) "Sample selection bias as a specification error," *Econometrica*, Vol. 47, pp. 153–161.

Heckman, James J. and Robert J. Willis (1977) "A beta-logistic model for the analysis of sequential labor force participation by married women," *Journal of Political Economy*, Vol. 85, pp. 27–58.

Hendren, Nathaniel (2017) "Knowledge of Future Job Loss and Implications for Unemployment Insurance," *American Economic Review*.

Hofferth, Sandra L., Sarah M. Flood, and Matthew Sobek (2017) "American Time Use Survey Data Extract Builder: Version 2.6," [dataset]. College Park, MD: University of Maryland and Minneapolis, MN: University of Minnesota. https://doi.org/10.18128/D060.V2.6.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow (2013) "The allocation of talent and us economic growth," NBER Working Paper.

Hyatt, Henry R. and James R. Spletzer (2017) "The recent decline of single quarter jobs," *Labour Economics*, Vol. 46, pp. 166–176.

Jones, Larry E., Rodolfo E. Manuelli, and Ellen R. McGrattan (2015) "Why are married women working so much?" *Journal of Demographic Economics*, Vol. 81, pp. 75–114.

Juhn, Chinhui (1992) "Decline of male labor market participation: The role of declining market opportunities," *The Quarterly Journal of Economics*, Vol. 107, pp. 79–121.

——— (2003) "Labor market dropouts and trends in the wages of black and white men," *Industrial & Labor Relations Review*, Vol. 56, pp. 643–662.

Juhn, Chinhui, Kevin M. Murphy, and Robert H. Topel (1991) "Why has the natural rate of unemployment increased over time?" *Brookings Papers on Economic Activity*, Vol. 1991, pp. 75–142.

——— (2002) "Current unemployment, historically contemplated," *Brookings Papers on Economic Activity*, Vol. 2002, pp. 79–116.

Kaplan, Greg (2012) "Moving back home: Insurance against labor market risk," *Journal of Political Economy*, Vol. 120, pp. 446–512.

Keynes, John Maynard (1930) "Economic possibilities for our grandchildren," *Essays in Persuasion*, pp. 358–73.

Knowles, John A. (2012) "Why are married men working so much? An aggregate analysis of intra-household bargaining and labour supply," *Review of Economic Studies*, Vol. 80, pp. 1055–1085.

Krueger, Alan B. (2017) "Where Have All the Workers Gone? An Inquiry into the Decline of the US Labor Force Participation Rate," *Brookings Papers on Economic Activity*.

Kydland, Finn E. and Edward C. Prescott (1982) "Time to build and aggregate fluctuations," *Econometrica: Journal of the Econometric Society*, pp. 1345–1370.

Ljungqvist, Lars and Thomas J. Sargent (2007) "Do taxes explain European employment? Indivisible labor, human capital, lotteries, and savings," *NBER Macroeconomics Annual*, Vol. 21, pp. 181–246.

Lucas, Robert E. (1977) "Understanding business cycles," in *Carnegie-Rochester conference series on public policy*, Vol. 5, pp. 7–29, Elsevier.

MaCurdy, Thomas, David Green, and Harry Paarsch (1990) "Assessing empirical approaches for analyzing taxes and labor supply," *Journal of Human Resources*, pp. 415–490.

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.

Moffitt, Robert A. (2012) "The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations," *Brookings Papers on Economic Activity*, pp. 201–264.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak (2011) "Internal migration in the United States," *The Journal of Economic Perspectives*, Vol. 25, pp. 173–196.

Mulligan, Casey B. (2001) "Aggregate implications of indivisible labor," *Advances in Macroeconomics*, Vol. 1.

——— (2002) "A century of labor-leisure distortions," NBER Working Paper.

Pencavel, John (1986) "Labor supply of men: a survey," *Handbook of Labor Economics*, Vol. 1, pp. 3–102.

——— (2002) "A cohort analysis of the association between work hours and wages among men," *Journal of Human Resources*, pp. 251–274.

Prescott, Edward C., Richard Rogerson, and Johanna Wallenius (2009) "Lifetime aggregate labor supply with endogenous workweek length," *Review of Economic Dynamics*, Vol. 12, pp. 23–36.

Rogerson, Richard (1988) "Indivisible labor, lotteries and equilibrium," *Journal of Monetary Economics*, Vol. 21, pp. 3–16.

Rogerson, Richard and Johanna Wallenius (2013) "Nonconvexities, retirement, and the elasticity of labor supply," *The American Economic Review*, Vol. 103, pp. 1445–1462.

Rubin, Donald B. (1976) "Inference and missing data," *Biometrika*, Vol. 63, pp. 581–592.

Schmitt, John (2003) "Creating a consistent hourly wage series from the Current Population Survey's Outgoing Rotation Group, 1979-2002," Working Paper.

Schmitt, John and Kris Warner (2011) "Ex-offenders and the Labor Market," *WorkingUSA*, Vol. 14, pp. 87–109.

Shenhav, Naama (2016) "Bribed to Wed? Family Formation and Labor Market Responses to the Gender Wage Gap," Working Paper.

Triest, Robert K. (1990) "The effect of income taxation on labor supply in the United States," *Journal of Human Resources*, pp. 491–516.

Winship, Scott (2017) "What's Behind Declining Male Labor Force Participation: Fewer Good Jobs or Fewer Men Seeking Them?", Mercatus Center.

Ziliak, James P. and Thomas J. Kniesner (1999) "Extimating Life Cycle labor Supply Tax Effects," *Journal of Political Economy*, Vol. 107, pp. 326–359.

A Additional Results

Robustness of In-and-Out Contribution

Table A.1: CPS In-and-Outs by True Non-participation Duration

	CPS In-and-Outs	CPS Dropouts
% Short-term (< 24 months)	81	6
% Long-term (< 24 months)	15	94
% Left-censored	4	0
% Right-censored	0	0

Source: SIPP 2008 panel. Individuals are classified as in-and-outs or dropouts based on the CPS window. Sample is limited to observations with at least 24 months padding on either side of the window. Table shows the share of non-participation among those classified as in-and-outs or dropouts that is part of non-participation spells less than 24 months in length, greater than 24 months in length, or unknown due to left- or right-censoring. All statistics are computed using survey weights.

Table A.2: Robustness of In-and-Outs' Contribution to Participation Decline since 1977

Year	(1) Baseline	(2) SIPP Adjusted	(3) De-NUNified	(4) Dropping Proxy Responses	(5) Including Want to Work
1980	0.3	0.3	0.3	0.3	0.3
1990	0.7	0.6	0.7	0.6	0.7
2000	1.3	1.2	1.3	1.2	1.0
2010	2.3	2.1	2.3	2.2	1.5
2015	2.2	2.0	2.1	2.2	2.3

Note: All units are percentage points. Baseline contributions to participation decline are computed using equation 1 applied to the CPS. Other columns make adjustments to account for measurement error as described in the text. All statistics are computed using survey weights.

Quality of Life

Table A.3: Differences in Quality of Life

	Dropouts	In-and-outs	Always Participators
Life Satisfaction (1-10)	5.9	6.3	7.1
% Took Pain Medication Yesterday	61	29	19
% With Any Physical or Cognitive Difficulty	58	24	8
% With Physical Difficulty	*40*	*8*	*1*
% With Mobility Difficulty	*26*	*4*	*0*
% With Difficulty Remembering	*21*	*5*	*0*
% Very Well Rested	27	42	38

Source: IPUMS CPS matched longitudinally, 2010-2013. IPUMS American Time Use Survey, 2010-2013. All statistics are computed using survey weights.

Geographic Distribution

Figure A.1: Changes in Participation Across States

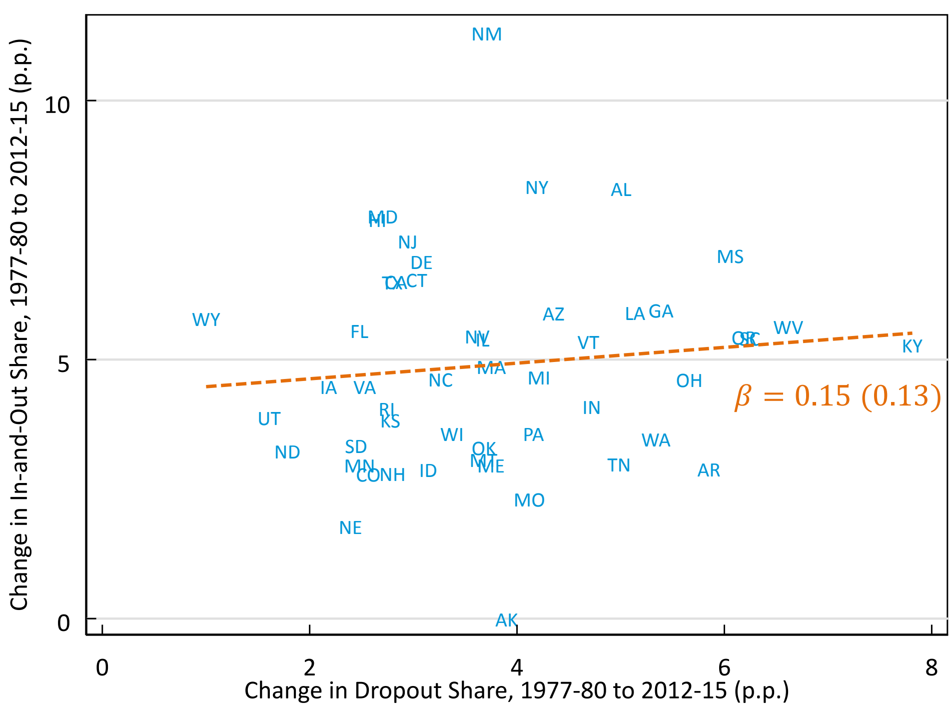

Source: IPUMS CPS matched longitudinally, 1977-2015. Slope is reported with heteroskedasticity-robust standard error in parentheses. All statistics are computed using survey weights.

Industries and Occupations

Table A.4: Rise of In-and-Outs Across Industries

Industry	ΔIn-and-Outs, 1977-2015 (p.p.)	Industry	ΔIn-and-Outs, 1977-2015 (p.p.)
Agriculture	6.5	Real Estate	0.2
Mining	1.6	Professional Services	3.6
Utilities	7.1	Administrative Services	0.3
Construction	3.0	Education	8.0
Manufacturing	3.6	Health	2.8
Wholesale Trade	3.3	Entertainment	1.1
Retail Trade	6.3	Food and Hospitality	1.0
Transportation	4.9	Other Services	4.5
Information	7.0	Public Administration	1.2
Finance and Insurance	1.6	Multiple or Unknown	4.5

Source: IPUMS CPS, 1977-2015. In-and-outs are defined as men ages 25-54 with between 1 and 7 months in the labor force out of 8 total. Individuals are categorized to industries based on the industry (or industries) they work in when employed. Individuals who are not employed in any of their 8 CPS responses are excluded. For each industry, I report the increase in the share of individuals within that industry who are in-and-outs between 1977 and 2015. All statistics are computed using survey weights.

Table A.5: Rise of In-and-Outs Across Occupations

Occupation Group	ΔIn-and-Outs, 1980-2015 (p.p.)	Occupation Group	ΔIn-and-Outs, 1980-2015 (p.p.)
Management, Professional, Technical	2.3	Production and Crafts	4.8
Admin. Support & Retail Sales	3.5	Machine Operators	3.7
Low-Skill Services	3.6	Transportation, Construction, Mining, Agriculture	3.7

Source: IPUMS CPS, 1980-2015. In-and-outs are defined as men ages 25-54 with between 1 and 7 months in the labor force out of 8 total. Individuals are categorized to occupations based on the occupation of their job when they are employed. Individuals who are not employed in any of their 8 CPS responses are excluded. For each occupational category, I report the increase in the share of individuals within that category who are in-and-outs between 1980 and 2015. All statistics are computed using survey weights.

Tax Rates

Figure A.2: Income Tax Rates of In-and-Outs

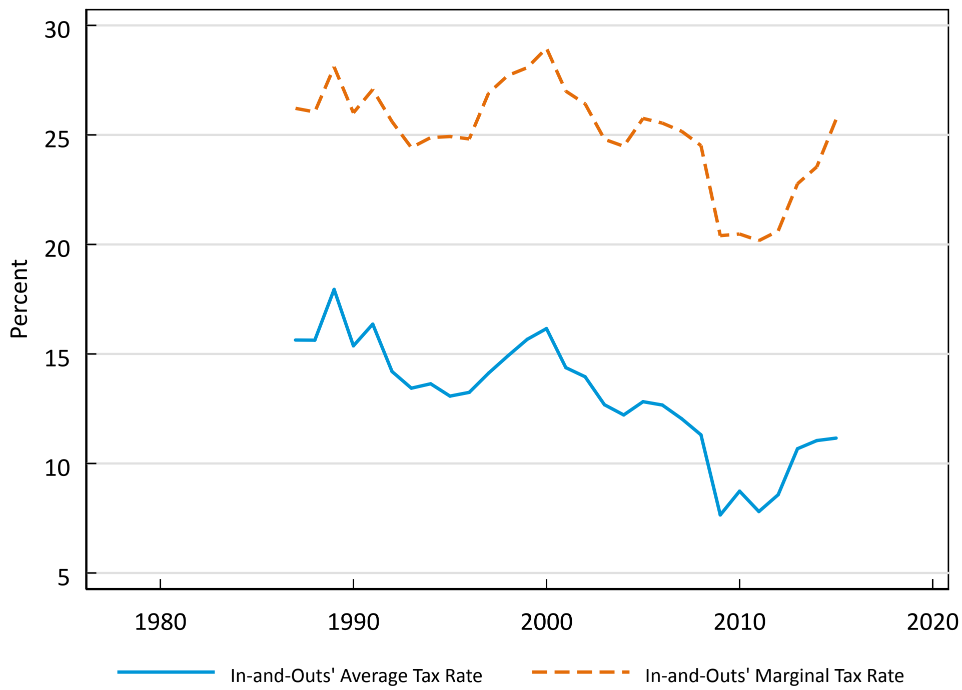

Source: IPUMS monthly CPS matched to IPUMS March CPS, 1989-2015. Tax rates are computed for all individuals in the sample using NBER's TAXSIM calculator (details in Appendix B). These individual-level tax rates are then averaged across all in-and-outs. All statistics are computed using survey weights.

B Datasets

This section describes several aspects of the datasets I construct for this analysis. First, I discuss the methods used to link CPS responses both longitudinally over time as well as between the monthly CPS and March CPS supplement. Next, I describe how I construct wage series along with several alternative imputation procedures. Finally, I outline how I estimate tax rates for individuals.

B.1 Linking Responses

I link CPS data in two ways. First, I link responses to the monthly CPS longitudinally over time. Second, I link responses from the monthly CPS in March to the March CPS supplement.

Figure A.3: Comparing Declining Participation Across Samples

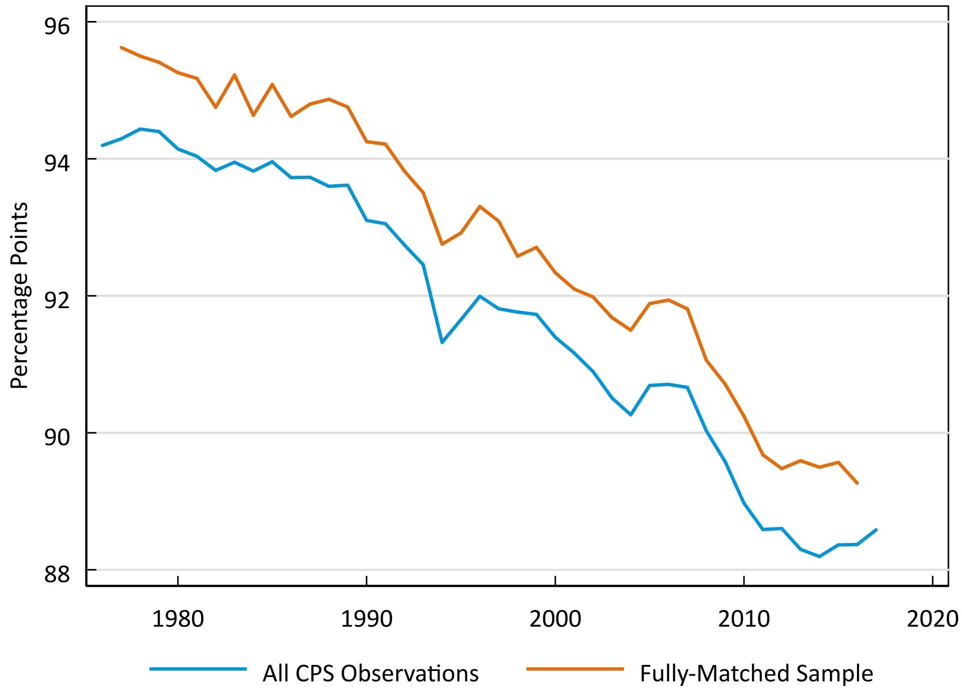

Source: IPUMS CPS 1976-2016, men ages 25-54. Fully-matched sample includes only individuals who can be matched across all eight responses. All statistics are computed using survey weights.

Longitudinal Linking The CPS has a 4-8-4 rotation group structure whereby individuals are interviewed for 4 consecutive months, rotate out for 8 months, and then are interviewed for another 4 months. As a result, individuals' responses can be matched to form a panel of 8 months covering a 16 month period. However, not all individuals can be successfully matched, since respondents are not followed if they move addresses or stop responding to the survey.

I use the method of Drew et al. (2014) to match individuals' responses across CPS interviews, as implemented in the IPUMS CPS Flood et al. (2017). This approach relies on mechanical matches between interviews based on longitudinal links provided in the CPS, as opposed to matching based on characteristics. In most years, about 60% of prime age men can be matched for all 8 responses, similar to the rates reported for the overall population by Drew et al. (2014).

Appendix Figure A.3 shows the evolution of the participation rate for all respondents as well as for those who can be matched to all 8 months. The fully-matched group is clearly a selected sample, as their participation rate is consistently higher than the average. However, the decline in participation is nearly identical for the two groups, so that although the fully-matched sample is selected, the degree of this selection has not changed over time.

Linking Monthly CPS to March CPS All respondents to the March monthly CPS are also administered a supplemental survey. However, since the data files for this supplement are released separately the responses must be matched to the monthly CPS to compare responses. Additionally, in some years the March CPS files are released with a different set of identifiers than the monthly CPS, making it impossible to match these responses.

To match responses between the two surveys, I utilize the method of Flood and Pacas (2016). This method focuses on mechanical matches using the identifiers released in the two surveys. I match responses during the 1989-2015 period, in which all monthly CPS responses can be matched to the corresponding March CPS response. However, not every March CPS response can be matched to a monthly CPS response, since the March CPS includes some individuals who do not appear in the monthly CPS.

B.2 Wages

I use several measures of wages in this analysis. In this section, I start by outlining the baseline measure of wages, which does not include those who report missing wages. Then I turn to different imputation procedures to estimate the wages available to those reporting missing wages.

Baseline Wage Measure To measure the wages of employed men, I use information from the CPS Outgoing Rotation Group (ORG) interviews. The ORG interviews take place for all individuals in their 4th and 8th month in the CPS sample and consistent of several questions about work, including nominal wages and hours worked.

From the ORG interviews, I construct a measure of real hourly wages. I start by dropping all observations with missing wages. Many individuals report wages at a weekly frequency instead of hourly. These weekly earnings are subject to a topcode for privacy reasons. I replace topcoded observations with the mean of weekly earnings above the topcode, assuming that weekly earnings follow a lognormal distribution (this method is outlined in Schmitt (2003)). After this adjustment, I convert weekly wages into hourly wages by dividing by hours worked per week. I convert nominal hourly wages to real wages using the PCE price index. Finally, I drop outlier observations, defined as hourly wages below $1/hour or above $300/hour in 2016$.

Imputations I use three different imputation procedures to estimate the wages available to those with missing wage information.

First, wages can be imputed non-parametrically using the method of Juhn (1992). In this approach, I estimate the distribution of wages within each year, specifically the average wage within each decile of the wage distribution. Each individual with missing wages is assigned probabilities of appearing in each decile based on the observed distribution of individuals with the same level of labor force attachment, where attachment is measured by the number of months in the labor force out of 8 in the CPS. In this way, each individual with missing wages is imputed the full distribution of wages based on similarly attached individuals.

I also use the regression imputation method of Blau and Kahn (2007). For each year, I regressed wages on a quadratic in age, 4 education dummies (less than high school, high school graduate, some college, and college graduate), 4 racial group dummies (white non-Hispanic, black, non-black Hispanic, and other), 3 marital status dummies (married, widowed/divorced, and never married), 9 census division dummies, a dummy for living in an urban area, and the total number of months the individual spent unemployed in the CPS (out of 8). I take the predicted values from this regression and use them as imputed wages for those with missing wage information.

Lastly, I impute wages using the predicted values from a regression with a selection correction as in Heckman (1979). I use the same covariates as in the regression imputation above, and use weekly spousal wages as the excluded variable to identify the selection equation. For any observation with missing wage information, I impute wages with the predicted value from the selection-corrected regression.

B.3 Tax Rates

I use the NBER TAXSIM program to estimate state and federal taxes for the sample of households I can match from the basic monthly CPS to the March supplement. To calculate taxes, I make the following assumptions: 1) I code all married couples as joint tax filing units and all other individuals as single filers; 2) I ignore mortgage interest, rent paid, child care expenses, and short-term capital gains (unavailable in March CPS).

I compute the marginal tax rate for an individual as the sum of the marginal federal income tax rate, the marginal state income tax rate, and the marginal payroll tax rate (including both employee and employer payroll taxes). The average tax rate is computed by taking total tax payments and dividing by total taxable income for each tax unit. To avoid problems with outliers, I winsorize both marginal and average tax rates at the 1% and 99% levels.

C Labor Supply of Men with Partners

This section outlines a model of household labor supply which shows how the income of other earners can affect individuals' labor supply. Using plausible parameter values as estimated by the literature, I find that the approximate magnitude of the wealth effect is very similar to the estimated wealth effect using variation in female wages across demographic groups from section 4.3.

The household's preferences over consumption c and labor provided by two adults n_1 and n_2 are given by $u(c, n_1, n_2)$ and the household is subject to a budget constraint, so that the household's problem can be expressed as

$$\max_{c,n_1,n_2} u(c, n_1, n_2) \quad s.t.\, c = w_1 n_1 + w_2 n_2 + Y$$

where w_1 and w_2 are post-tax wages for the two adults and Y includes all unearned income. For simplicity, I assume that changes in labor supply come from the in-and-out margin alone, although this is not a very restrictive assumption since in an augmented model with additional margins of labor supply the in-and-out margin is likely to account for all changes in labor supply anyway (Prescott et al., 2009; Rogerson and Wallenius, 2013). The first order conditions from optimization equate the marginal rate of substitution with the returns to work for each adult[51]

$$\frac{-u_{n_1}}{u_c} = w_1; \quad \frac{-u_{n_2}}{u_c} = w_2$$

For simplicity, I make an assumption that the disutility of labor is separable, i.e. $\frac{\partial^2 u}{\partial n_1 \partial n_2} = 0$.[52] Log-linearizing these conditions and re-arranging gives

$$\tilde{n}_1 = \theta_1 \tilde{w}_1 + \phi_1 \alpha_2 (\tilde{w}_2 + \tilde{n}_2) + \phi_1 \alpha_Y \tilde{Y} \tag{11}$$
$$\tilde{n}_2 = \theta_2 \tilde{w}_2 + \phi_2 \alpha_1 (\tilde{w}_1 + \tilde{n}_1) + \phi_2 \alpha_Y \tilde{Y} \tag{12}$$

where $\tilde{x}$ represents the percentage change in the variable x, θ_i is the uncompensated labor supply elasticity for adult i, ϕ_i is the income elasticity of labor supply, and α_i is the budget share of income earned by adult i. Importantly, changes in labor supply $\tilde{n}_1$ and

[51] I focus here on the case where both adults are at an interior solution for labor supply. For men, this is a trivial assumption since my empirical analysis examines labor supply among households containing men who are in the labor force at least some of the time. For women, this is a more binding assumption, since some women in these households may be permanently out of the labor force and thus the relevant first order condition for these individuals would be an inequality rather than equality.

[52] When examining the labor supply of household members using the same event study analysis as in sections 4.4 and 4.5, I find little change in the labor supply of other household members during the period when in-and-outs are out of the labor force. This suggests that leisure complementarities are minimal in the short-run.

$\tilde{n}_2$ are jointly determined in response to changes in wages.

Consider a change in the wages available to the second adult. In response, labor supply of both adults will adjust to a new equilibrium level. The equilibrium response of each adults' labor supply can be decomposed into several components, including a wealth effect on the labor supply of the first adult. Letting $\tilde{n}_1^*$ and $\tilde{n}_2^*$ denote the equilibrium response of each adults' labor supply in response to a change in wages $\tilde{w}_2$, holding all else constant, equations 11 and 12 above can be rewritten as:

$$\tilde{n}_1^* = \underbrace{\phi_1\alpha_2(1+\theta_2)\tilde{w}_2}_{\text{Wealth Effect}} + \underbrace{\phi_1\alpha_2(\tilde{n}_2^* - \theta_2\tilde{w}_2)}_{\text{Remainder}} \tag{13}$$

$$\tilde{n}_2^* = \underbrace{\theta_2\tilde{w}_2}_{\text{Direct Effect}} + \underbrace{\phi_2\alpha_1\tilde{n}_1^*}_{\text{Remainder}} \tag{14}$$

The change in the first adults' labor supply in response to an increase in the wages of the second adult is mostly governed by the wealth effect, since the remaining terms are proportional to the product of income elasticities and therefore are likely to be small.
Proposition. *If the product of income elasticities and budget shares is small, i.e.* $\phi_1\phi_2\alpha_1\alpha_2 \ll 1$, *then the equilibrium response of the first adult's labor supply* $\tilde{n}_1^*$ *to a change in wages* $\tilde{w}_2$ *will be approximately equal to* $\phi_1\alpha_2(1+\theta_2)\tilde{w}_2$.

The magnitude of the wealth effect can be approximated using estimates of these parameters from the literature. I take the average income elasticity for men across six studies (Eissa and Hoynes, 2004; MaCurdy et al., 1990; Moffitt, 2012; Pencavel, 1986; Triest, 1990; Ziliak and Kniesner, 1999), yielding $\phi_m = -0.06$. The female uncompensated labor supply elasticity is estimated by Eissa and Hoynes (2004) to be $\theta_f = 0.27$. For the budget share of female labor income, I use a sample of March CPS households containing prime age men with a spouse or unmarried partner and compute the median share to be $\alpha_f = 0.3$. Putting these together yields an estimated wealth effect of -0.02, nearly identical to the estimate computed from variation in female wages across demographic groups in section 4.3.

D Time Use on Video Games

This section reconciles the results on time use presented in Table 5 with previous results from the literature. Aguiar et al. (2017) point out that non-employed young men spend a growing fraction of their time playing video games. They interpret this fact through a model of time use as implying that the quality of video games has improved, which they then show can account for some of the decline in labor supply among this population since 2000. However, I find that in-and-outs spend relatively little time on video games

Table A.6: Video Game Time Crosswalk

	2004-07	2012-15
Sample as in Table 5	1.6	
Breaking Out 2004-07 and 2012-15	1.7	1.0
Including All Matched Non-Participants	2.9	1.4
Including All Non-Participants	3.2	2.1
Including All Non-Employed	2.7	2.5
Using Whole Week Time Use	2.4	2.6
Changing to Men Ages 21-30	3.2	6.5
Excluding Full-Time Students	3.5	6.5
Reported by Aguiar et al. (2017)	3.4	5.9

Notes: Top row is estimated as in Table 5. Each subsequent row is one step on the crosswalk to the Aguiar et al. (2017) estimate, which is reported in the last row. First, the pooled average is broken out into the 2004-07 and 2012-15 averages. Then non-in-and-out non-participants who can be matched to the monthly CPS are added, followed by all non-matched non-participants and then all non-employed. Next, I include time use on weekend days, as opposed to weekdays only. After this, I change the age restriction from 25-54 to 21-30. Finally, I drop full-time students ages 24 and under. This second-to-last line is the same concept as Aguiar et al. (2017), but may not be identical if there are differences in weighting. All statistics are computed using survey weights.

and experience little increase in time spent on video games when leaving the labor force.

There are many differences between these two analyses that could account for the apparent discrepancy. While I focus on in-and-outs, Aguiar et al. (2017) look at non-employment overall. Additionally, they focus on young men between the ages of 21 and 30, while I examine men between the ages of 25 and 54. There are also some differences in time period and sample restrictions. To examine the influence of these differences, I recompute the time spent on video games, changing one aspect at a time to fully crosswalk the two sets of estimates.

Appendix Table A.6 shows how the estimated hours per week spent on video games changes with each adjustment. The 0.2 hours per day reported in Table 5 is equivalent to 1.6 hours per week. Breaking out separate time periods, as is done by Aguiar et al. (2017), doesn't substantially alter the result. However, including dropouts and other non-employed individuals raises the average time spent on video games substantially, approximately double the average of just in-and-outs within each time period. This implies that these other groups spend significantly more time on video games than in-and-outs do. Expanding time use to include weekend days makes little difference. Changing the sample to look at 21-30 year olds increases the estimated time spent on video games in both time periods, but especially in 2012-2015. Dropping full-time students has little effect. The second-to-last row gives estimates using the same concept as Aguiar et al. (2017), but although the results are close they are not identical for unknown reasons.

The two deviations between the analyses that make the most difference are: 1) the choice of in-and-outs versus all non-employed, and 2) the focus on 25-54 year olds instead of 21-30 year olds. Both of these choices appear to contribute about equally to the difference in estimates.

E Evolution of In-and-Outs' Income

One potential problem with the event study analysis of section 4.5 is that the sample of in-and-outs is defined using ex-post information, specifically whether individuals returned to the labor force after a short period of time. In this way, the analysis may exclude individuals who anticipated taking a short break out of the labor force, but ended up experiencing large declines in available wages and remained out of the labor force. This would suggest that while in-and-outs do not suffer permanent income loss, this is merely because they were lucky rather than because a short break out of the labor force cannot incur a permanent cost.

The extent of this problem can be gauged by using a sample selected based on ex-ante characteristics only, since this avoids using ex-post information. Using the SIPP sample, I examine all employment to non-participation separations and use pre-separation characteristics to predict the duration of the non-participation spell. This relies on the notion that individuals with similar characteristics are likely to spend similar amounts of time out of the labor force, or put another way that in-and-outs can be separated from dropouts by only examining their pre-separation characteristics. This is an empirically testable statement. I use all individuals who transition from $E \rightarrow N$ at least two years before the end of the sample and use as covariates their demographic, household, job, and income characteristics from 4 months before the separation.[53]

To predict the duration of non-participation based on pre-separation characteristics, I use a machine learning algorithm known as Gradient Boosted Trees (GBT).[54] GBT fits the outcome by iteratively applying a series of small decision trees. The first tree is fit to the raw data, splitting branches of the tree to maximize the share of variation explained by the tree. The next tree is fit to the residuals from the first tree and the process repeats, with each tree in the series fitting the residuals of the previous one. However, each tree is weighted according to the marginal improvement in fit, resulting in later

[53] The restriction that the transition must occur at least two years before the end of the sample avoids including transitions where the non-participation spell duration is right-censored under two years. It is unclear whether these excluded observations are in-and-outs or dropouts due to this censoring. I use covariates from four months before the separation so that all characteristics are measured in a previous SIPP wave, since characteristics may be mechanically correlated within a SIPP wave.

[54] GBT is described in detail in Hastie, Tibshirani and Friedman (2009). I use the LightGBM package to implement this algorithm.

Figure A.4: Importance of SIPP Characteristics in Predicting Non-Participation Duration

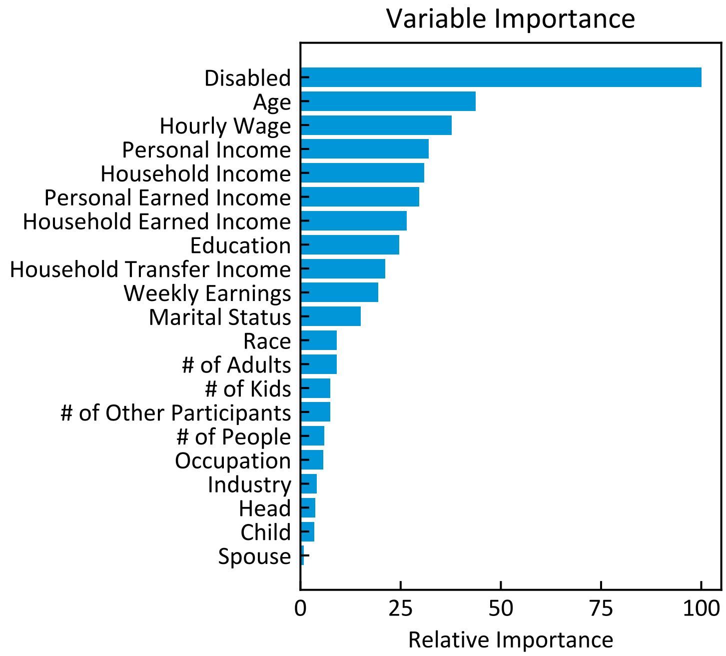

trees receiving less weight as the algorithm reaches a point of diminishing returns in improving the fit of the model. Each tree is constrained to be very shallow. By repeatedly applying shallow trees to the residuals of past iterations, this method can approximate very flexible functional forms, with later trees correcting for misspecifications introduced by earlier trees. I use a total of 50 trees with a learning rate of 0.05, as selected by cross-validation.

Each characteristic's contribution to the prediction can be quantified. I compute the total "gain" provided by a variable by summing up the reduction in the sum of squared residuals from every time that variable is used to split a branch across all of the 50 trees. Appendix Figure A.4 shows this measure for each variable, expressed as a share of the highest gain. Disability is the most important characteristic for separating in-and-outs from dropouts, accounting for more than double the importance of any other variable. Age, income, and education are also judged to be important predictors of time spent out of the labor force.

Next, I examine whether this algorithm is able to adequately separate in-and-outs from dropouts. I order individuals by their predicted duration and categorize all individuals with predicted duration below some threshold to be predicted in-and-outs. For a given threshold, this results in some true positive rate, i.e. the share of actual in-and-outs who are categorized as such, and a false positive rate, i.e. the share of dropouts who are incorrectly categorized as in-and-outs. I vary this threshold to trace out a curve trading off these two different rates, known in the machine learning literature as a Receiver-

Figure A.5: ROC Curve for Predicted In-and-Outs

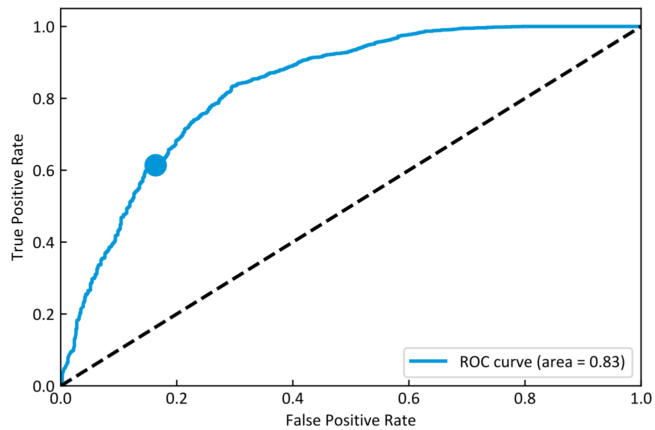

Operating Characteristic (ROC) curve, which is plotted in Appendix Figure A.5. The ROC curve can be compared to the 45º degree line to measure the predictability of in-and-outs. If in-and-outs are completely random, the ROC curve will coincide with the 45º line, while if in-and-outs are perfectly predictable, the ROC curve will be equal to 1 everywhere. In this case, the ROC curve is quite a bit above the 45º line, indicating that in-and-outs can be fairly easily separated from dropouts. The area under the ROC curve measures the deviation from randomness and in this case it is equal to 0.83, indicating a strong degree of prediction, although not quite complete prediction.

Using these predictions, I construct a sample of predicted in-and-outs and repeat the event study from above. I select a sample of predicted in-and-outs that captures about 60% of actual in-and-outs while including less than 20% of actual dropouts (marked with a dot in Appendix Figure A.5). As above, I estimate the following equation for an outcome Y,

$$\Delta Y_{i,t,t+k} = \beta^{(k)} \mathbf{1}\left(\text{In-and-Out}_{i,t}\right) + \delta X_{i,t} + \Delta\epsilon_{i,t,t+k} \tag{15}$$

varying k between -6 and 24 to cover a two-and-a-half year window, controlling for a cubic in age as well as time fixed effects.

First, I show that the predicted in-and-outs behave like actual in-and-outs using this event study. I estimate equation 15 using labor force participation as the outcome and plot the resulting coefficients in Appendix Figure A.6. Predicted in-and-outs return to the labor force quite quickly, with the majority returning within the first six months. By twelve months after the separation, predicted in-and-outs have similar participation rates as actual in-and-outs. Predicted dropouts, on the other hand, remain disconnected from the labor force even after two years, as do actual dropouts. The event study shows how the choice of trading off true positive rates and false positive rates affects the groups, as I have chosen a sample of predicted in-and-outs that contain very few actual dropouts,

Figure A.6: Evolution of Predicted In-and-Outs' Participation

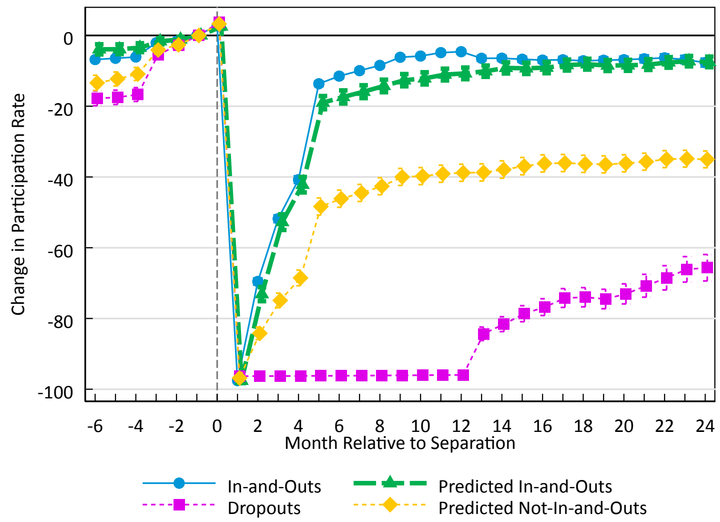

Sample: 25-54 year old men in the SIPP. In-and-Outs are employed in month 0, nonparticipants in month 1, and are employed again by at least month 12. Dropouts are employed in month 0 and nonparticipants in months 1-12. For predicted in-and-outs and dropouts, re-employment time is predicted from covariates four months before the separation using gradient boosted trees and individuals are categorized based on whether the predicted time is less than a threshold delivering a false positive rate below 20%.

but exclude some actual in-and-outs, resulting in a gap between predicted dropouts and actual dropouts. However, I examine the evolution of income for the predicted in-and-outs only and not for predicted dropouts, so this tradeoff is sensible.

Next, I turn to the evolution of income for the group of predicted dropouts. Appendix Figure A.7 shows the estimated evolution of income for predicted in-and-outs, actual in-and-outs, and unemployed job losers, where all three sets of estimates are computed using equation 15. While predicted in-and-outs experience an even larger decline in personal income than actual in-and-outs at the time of separation, this gap narrows as both groups return to the labor force and the income of the two groups is statistically indistinguishable by one year after the separation. Both groups experience no significant change in income between the month before the separation and two years after the separation, while unemployed job losers experience significantly lower income.

Figure A.7: Evolution of Predicted In-and-Outs' Income

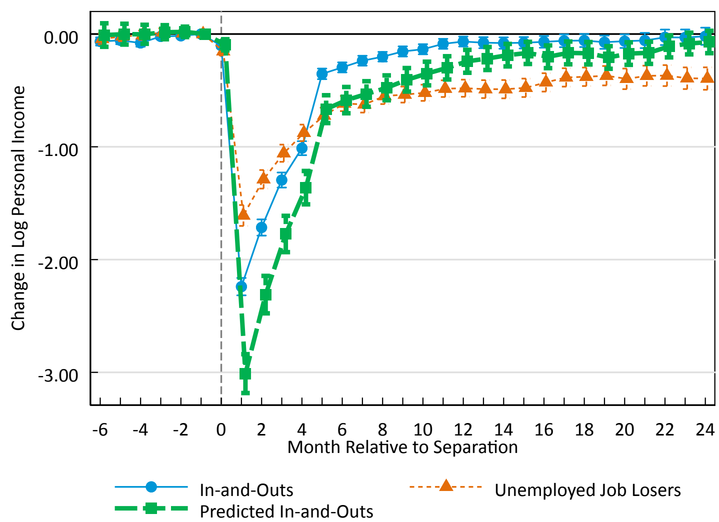

Sample: 25-54 year old men in the SIPP. In-and-Outs are employed in month 0, nonparticipants in month 1, and are employed again by at least month 12. Unemployed job losers are employed in month 0 and fired or laid off but looking for work in month 1. Re-employment time is predicted from covariates four months before the separation using gradient boosted trees and individuals are categorized based on whether the predicted time is less than a threshold delivering a false positive rate below 20%.