What Quits and Layoffs Reveal About the Business Cycle

Kathrin Ellieroth
Colby College

Amanda Michaud *
Federal Reserve Bank of Minneapolis

Preliminary Draft

This Version: June 23, 2026

Please click here for latest version: Latest version

First Version: April 2024

Abstract

We introduce monthly CPS series that classify every separation by reason (quit or layoff) and destination (unemployment or non-participation), and use them to confront the workhorse model of labor market flows. The standard model overproduces quits, sends laid-off workers to unemployment when a third of them leave the labor force, and gets the cyclical direction of labor force attachment backwards. Adding selective layoffs and random quits repairs these facts and overturns what labor supply does over the cycle: a recession is partly stabilized from within. Two channels buffer employment and amplify unemployment: (i) labor supply as marginal workers hoard their jobs and the displaced keep searching, and (ii) selection as layoffs shift from marginal toward attached workers. These mechanisms reorganize the drivers of the cycle and sharpen common questions: recessions feature near zero changes in TFP and welfare costs 59% lower than the standard model.

JEL CODES: E32, J2, J6

1 Introduction

The workhorse macroeconomic model of gross worker flows interprets the labor market through two channels: flows from employment to unemployment are primarily layoffs, and flows from employment to non-participation are primarily quits. Layoffs are random, a separation shock that hits all workers with equal probability, and quits are choices made by workers whose value of working has drifted

*E-mail: kathrin.ellieroth@colby.edu or amanda.michaud@gmail.com. This work implements methodology developed in training at the 2018 Inaugural Summer Data Workshop: Using the Panel Component of IPUMS CPS hosted by the Minnesota Population Center. The data we construct in this paper will be available at https://sites.google.com/qlmonthly.com/home and are updated monthly.

We thank Justin Bloesch, Domenico Ferraro, Jonathan Heathcote, Zhen Huo, Andre Kurmann, Simon Mongey, Lukas Nord, Xincheng Qiu, Andrea Raffo, Victor Rios-Rull, Richard Rogerson, Abigail Wozniak; participants at FRB Minneapolis, LAEF, NBER, SED, Midwest Macro, SAMM, and other seminars and conferences; and our active community of data users for the discussions and helpful feedback. Sara Canilang provided excellent research assistance. The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of Minneapolis, the Federal Reserve System, or the Board of Governors.

below their value of staying home. This structure, in the tradition of Krusell et al. (2017), underlies most quantitative work on the flow approach to business cycle analysis, and it underlies how researchers and policymakers read the data: unemployment measures displaced workers and non-participation measures changes in labor supply.

We construct new monthly series from the CPS, from 1978 to 2025, that classify every separation into non-employment by its *reason* (quit or layoff) as well as its *destination* (unemployment or non-participation). These series confront the standard model with facts it was never asked to match, and it fails on three of them. First, it does not produce the composition of separations: it generates too many quits, and, most starkly, it sends nearly every laid-off worker to unemployment, putting the share of layoffs flowing directly to non-participation at 22% against 35.5% in the data. Second, it has almost no flows from non-participation into search: its NU rate is 0.2% against 3.5% in the data, because too few of its non-participants are near the margin. Third, it misses the direction in which labor force attachment moves over the cycle. In the data, recessions make the jobless *more* attached: quits fall by a third, exits from unemployment to non-participation fall (the UN rate declines from 18% to 15%), and flows from non-participation into the labor force rise (the NU rate increases from 3.5% to 5.6%). The standard model generates muted quit and attachment responses and, as we show, assigns the resulting labor supply response the wrong sign. These are not small calibration misses. They are structural: the failures persist because the model has no mechanism that could repair them.

The diagnosis points to the repair. The standard model fails because all of its layoffs are random and all of its quits are marginal labor supply choices. The data require the opposite on both margins. Layoffs must be partly *selective*: firms shedding workers cut those with the least to lose, low surplus workers near the participation margin, and these workers, having just revealed themselves to be near the margin, often exit the labor force after displacement. And quits must be partly *random*: a large share of employment-to-non-participation flows reflect events far from the margin of participation (disability, family circumstances, a spouse's job) that no business-cycle-frequency change in wages or benefits will reverse. We add both ingredients to an otherwise standard Aiyagari economy with three labor market states. A continuation-cost shock generates endogenous selective layoffs of low-surplus workers via their decision rules, and a participation-constraint shock generates exogenous quits of workers who are firmly detached. We estimate the model by simulated method of moments targeting the levels and volatility of our new separation series in addition to the standard stocks and flows. With the two added ingredients, the model matches the composition of separations and reproduces rising attachment in recessions: quits fall, UN falls, NU rises, and share of layoffs to N fall all as in the data.

Getting the composition of separations right is not bookkeeping; it changes the answers to central questions about the business cycle. We organize the paper around four of them.

What does labor supply do over the cycle? It buffers recessions. In recessions, marginal workers hoard their jobs: quits fall by almost half a percentage point, the displaced stay more attached, and non-participants flow into search (the added-worker flow NU more than doubles). The net effect of these labor supply responses is to make employment fall *less* (by 1.0 percentage point at the trough) and measured unemployment rise *more* (by 0.3 percentage points): labor supply is an employment stabilizer and an unemployment amplifier. The standard model does not merely fail to reproduce this, it gets the sign wrong. It requires a larger decline in TFP, which leads to a deeper decline in employment through rising quits. Getting the composition of separations and the attachment flows right is what makes labor supply an employment stabilizer. Further, our analysis clarifies the mechanism of counter-cyclical labor supply: hoarding is concentrated in a thin slice of marginal workers and driven by the overall decline in job finding rates, while marginal non-participants supply the added-worker entry in response to changes in differential job finding rates in U versus N.

How cyclical is selection in separations? Selection is significant. In normal times 66% of quits and 31% of layoffs are selective, initiated by low-surplus workers near the margin. In recessions both shares collapse, to 19% and 18%: quits and layoffs become more random, reaching up the surplus distribution into the firmly attached as low-surplus marginal workers cling to their jobs. This cyclical de-selection accounts for most of the observed fall in the share of layoffs flowing to non-participation, and it works alongside labor supply: holding the composition of layoffs at its normal-times incidence, employment would fall 0.72 percentage points more. Moreover, we show this pattern is quantitatively consistent with Mueller (2017), who shows the pool of separators shifts strongly toward high-productivity workers in recessions.

What do business cycles cost, and what is UI worth? The welfare cost of a recession episode averages 0.63% of lifetime consumption in our model in comparison to 1.55% in the standard model. The difference comes largely from selection in two senses. The first is how the estimation infers the driver of business cycles: changes in selection in our model account for the fall in output per worker without needing a decline in TFP, a very costly feature of the standard calibration, in a welfare sense. The second is a genuine economic mechanism within our calibrated model: the average cost would be 1.0% in our model if layoffs were not selective. The two models also disagree about policy. Extending UI benefits in recessions raises the unemployment rate by 0.28 percentage points in our model (about a ninth of the total increase) by holding displaced marginal workers in measured unemployment, and it is the force that makes UN flows fall as in the data. The average worker would pay 0.12% of consumption to remove the extension, while the standard model scores it as approximately free.

Is unemployment a consistent measure of slack? The majority of flows to employment from non-employment originate in non-participation and this share changes over the business cycle suggesting that unemployment is an inconsistent measure of slack. The model, by contrast, prices every non-employed worker's willingness to work: we measure the jobless who want to work (WTW) as those whose value of employment exceeds the value of their current non-employed state. This stock is six times the unemployment stock in normal times and a third as cyclical in the GFC. The wedge between the two is policy and labor supply: in the GFC the UI extension alone moves the unemployment rate by 0.3 ppt while moving WTW by a tenth of a point. The unemployment rate is a measure of who is *searching*, and search responds to benefits; WTW is a measure of who would *work*, and that is the object that matters for slack. While our baseline calibration is for the GFC, we show that the divergence can be very different in other episodes, thus necessitating updated tracking of WTW which is exactly what our flow series and framework provide: the participation choices of laid-off workers identify, in close to real time, where the margin of the labor market sits.

We close by re-estimating the model on two further episodes. The 1981–82 recession is displacement in nearly pure form: job loss alone accounts for two-thirds of the unemployment rise, the era's less generous UI contributes far less than the Great Recession's extension does, and the selective share of layoffs falls sharply. The 2022–23 tight labor market is the mirror image of a recession on the aggregates but quantitatively exhibits a substantial decline in labor supply which held back employment growth and caused the small decline in the unemployment rate to substantially understate how tight the labor market had become.

Literature Review Our paper contributes to the empirical and theoretical literature that examine labor market flows between employment, unemployment, and non-participation.

Labor Market Flows and Transitions. A rich literature following work by Abowd and Zellner (1985) and Blanchard and Diamond (1990) has analyzed gross flows and transition rates between the labor market states. This body of research aims to understand the evolution of labor market flows across time, cross-sections, and business cycles. Shimer (2012) utilizes flow data on employment-unemployment transitions to construct job-finding and job-loss probabilities, assessing their relative importance to unemployment rate fluctuations. Similarly, Elsby et al. (2015) and Elsby et al. (2019) employ data on flows between employment, unemployment, and non-participation to analyze the contribution of each flow to labor force participation rate fluctuations. Others, such as Garibaldi and Wasmer (2005), Krusell et al. (2017), Cairó et al. (2022), Qiu (2022), and many more have used gross flows data to inform macroeconomic models of labor markets. Our contribution to this literature is to discipline these models with the composition of separations by reason and destination, and to show that the discipline is binding: the workhorse model cannot match it without selective layoffs and random quits.

Closest on the data side is Simmons (2023), who shows in SIPP and the UK LFS that allowing for heterogeneity in the reason for separation raises the measured contribution of separations to the ins of unemployment via a Shimer-style flow accounting. We reach a complementary conclusion using a structural model, and add the destination margin.

Marginal Workers and Labor Force Attachment. Our focus on marginal workers connects to literature examining labor force attachment. Jones and Riddell (1999) and Jones and Riddell (2006) identify a group of non-employed individuals who report wanting work but are not actively searching. This "waiting" group has transition rates to employment that are higher than other non-participants but lower than the unemployed.[1] More recently, Morchio (2020), Gregory et al. (2025), Ahn et al. (2023), identify marginal groups as individuals who flow often between states. We complement this work through a different identification approach, using quits and exits after layoff, to track this distribution closer to real time instead of relying on work histories.

Our analysis of marginal workers also relates to Barnichon and Figura (2015), who demonstrate that compositional changes in the unemployed population—including shifts in labor force attachment—substantially affect aggregate matching efficiency. Similarly, Hornstein et al. (2014) develop a framework where non-participants' heterogeneous labor force attachment influences aggregate job-finding rates.

Selection in Quits and Layoffs. Gibbons and Katz (1991) offer empirical evidence for selection in layoffs, showing that workers displaced by plant closings face shorter unemployment durations and smaller wage losses than those laid off for other reasons. More recently, Mueller (2017) finds that laid-off workers have substantially lower pre-displacement wages than those who quit, supporting the selection hypothesis. Chodorow-Reich and Karabarbounis (2016) document cyclical patterns in the opportunity cost of employment, and Mui and Schoefer (2021) show that selection into unemployment varies systematically over the business cycle. Our model nests both selective and random separations on each margin and lets the data determine the mix, in each phase of the cycle.

Unemployment Dynamics and Slack. Our work contributes to understanding unemployment movements not directly related to labor market slack. Ahn and Hamilton (2021) decompose unemployment fluctuations into flows related to entry, exit, and duration, finding significant variation unrelated to aggregate demand.

We contribute to the "unemployment volatility puzzle" literature pioneered by Shimer (2005) and further explored by Hagedorn and Manovskii (2008), Ljungqvist and Sargent (2017), and Mitman and Rabinovich (2019). We contribute by showing that layoffs are more frequent and less cyclically volatile than flows from employment to unemployment, and that labor supply appears to be countercyclical on the margin.

Mukoyama et al. (2018) document countercyclical search intensity, suggesting that labor sup-

[1] Flinn and Heckman (1983) test whether unemployment and non-participation are behaviorally distinct labor force states and find that unemployed workers have significantly higher job-finding rates than non-participants. Building on this, Krueger et al. (2014) examine the long-term unemployed and find that many become marginally attached to the labor force, cycling between active search and non-participation.

ply behavior—rather than just demand factors—shapes unemployment fluctuations. The role of unemployment insurance in distorting unemployment measures of slack has been investigated by Rothstein (2011) and Farber and Valletta (2015), who find that UI extensions modestly increase unemployment duration but primarily through reduced labor force exits rather than reduced job finding. This aligns with our finding that marginal workers are more likely to report as unemployed when benefits are more generous.

Cyclical Labor Supply. Most directly related to our central finding is the contemporaneous work of Graves et al. (2023), who estimate the response of labor-market flows to high-frequency monetary-policy shocks. They find, as we do, that labor supply cushions the employment decline through the response of quits to a falling job-finding rate, in a model in the tradition of Krusell et al. (2017).[2] We, however, obtain estimates about the role of labor supply that are less than half as large: the employment decline would be about a third larger without the labor-supply response in our model, against roughly 80% larger in theirs. The difference turns on how each model interprets the central fact that nearly 40% of laid-off workers exit the labor force. Because Graves et al. (2023) allow only random layoffs, their model must read this fact as evidence that 40% of *all* workers are marginal: they sit close enough to the participation margin to exit upon displacement, and it also attributes the entire cyclical fall in exits-after-layoff to a pure labor supply response. That delivers a very high labor-supply elasticity. Our model instead splits the fact between labor supply and selection: layoffs fall disproportionately on low-surplus workers, who are the ones likely to exit, and part of the cyclical movement is a shift in *who* is laid off rather than a change in workers' participation choices. Critically, our mechanism is validated by estimates of cyclical selection by Mueller (2017), which our model succeeds in replicating despite not targeting in the estimation.

2 Data and Methodology

2.1 Data source

We use monthly data from the Current Population Survey (CPS) from January 1978 to July 2024. The CPS is a rotating panel survey of approximately 60,000 households, conducted by the US Bureau of Labor Statistics. While primarily designed for cross-sectional analysis, the CPS's rotating panel structure allows us to match individuals across consecutive months, enabling the

[2] Our CPS quit and layoff series were developed independently, applying standard and publicly documented procedures for the panel component of the CPS: the survey's long-standing reason-for-separation questions for the unemployed (harmonized in IPUMS, Flood et al. (2023)) and for non-participants (collected since 1967 and long analyzed in the literature, e.g. Flaim (1969), Flaim (1973), Schwab (1974), Bednarzik and Klein (1977), and Gellner (1975)), with longitudinal matching following Madrian and Lefgren (1999). The series we construct were the first of their kind made publicly available; they classify every separation by reason and destination and are updated monthly.

From	To		
	E	U	N
E	f_{EE}	f_{EU}	f_{EN}
U	f_{UE}	f_{UU}	f_{UN}
N	f_{NE}	f_{NU}	f_{NN}

Table 1: Standard approach of flow rates in the CPS

computation of month-to-month labor market transitions.

2.2 Methodology

We classify flows from employment to both unemployment and non-participation by reason of separation. The goal is to newly classify four distinct flows:

- Employment to unemployment due to a quit (EUQ)
- Employment to unemployment due to a layoff (EUL)
- Employment to non-participation due to a quit (ENQ)
- Employment to non-participation due to a layoff (ENL)

The CPS short panel follows a 4-8-4 structure which allows us to observe individuals for 4 continuous months, followed by an 8-month break, and then another 4-month period. Due to the option of observing individuals for two consecutive months, researchers have frequently used the CPS to compute gross flows and transition rates (Abowd and Zellner (1985), Shimer (2012), Elsby et al. (2015), and many others). Most commonly, researchers have computed flow rates between the three labor market states employment (E), unemployment (U), and non-participation (N) to create a matrix of nine flow rates as shown in Table 1. The flows have been used to understand fluctuations in job finding and job loss rates, or to study the evolution of stocks such as the unemployment rate or employment-population ratio using a stock-flow analysis.

The standard approach often interprets flows between employment and unemployment as layoffs and flows between employment and non-participation as quits. We show this convention is not accurate. Flows into both unemployment and non-participation consist of both layoffs and quits.

We follow the standard methodology of computing gross flows with an important difference: we compute flow rates from employment to both unemployment and non-participation *by reason of separation*. Thus, we not only get employment to unemployment (EU) and employment to non-participation (EN) rates, but also employment to unemployment due to a quit (EUQ), employment to unemployment due to a layoff (EUL), employment to non-participation due to a quit (ENQ), and employment to non-participation due to a layoff (ENL), such that

$$f_{EU} = f_{EUQ} + f_{EUL} \tag{1}$$

$$f_{EN} = f_{ENQ} + f_{ENL} \tag{2}$$

Few papers distinguish separations into non-participation by quits and layoffs. Table 2 shows the contribution of this distinction to the standard approach of using the CPS to calculate flows. While this seems like a minor change, it allows researchers to use this data in important ways, such as (i) analyzing what fraction into unemployment and non-participation is due to a layoff vs. a quit; and importantly (ii) accurately observing total quits and total layoffs into non-employment, i.e

$$\text{Quits} = f_{EUQ} + f_{ENQ} \tag{3}$$

$$\text{Layoffs} = f_{EUL} + f_{ENL} \tag{4}$$

From	To		
	E	U	N
E	f_{EE}	$f_{EUQ} + f_{EUL}$	$f_{ENQ} + f_{ENL}$
U	f_{UE}	f_{UU}	f_{UN}
N	f_{NE}	f_{NU}	f_{NN}

Table 2: Our contribution to the standard approach

2.3 Decomposition into Layoffs and Quits

2.3.1 Unemployment

We are going to keep this section brief, since the distinction of a layoff or quit into unemployment in the CPS has been used in previous literature. In CPS IPUMS (Flood et al. (2023)), the variable to classify a separation into unemployment as a quit, layoff, or other is readily available and harmonized for all sample months. The survey asks all unemployed individuals why they became unemployed and distinguishes between workers who had lost jobs (due to temporary layoff, involuntary job loss, or ending of a temporary job), those who had quit jobs, those who were re-entering the labor force after an extended absence from the work force, and those who were seeking their first jobs (new entrants). We use these answers and classify a separation into unemployment as a layoff or quit as follows:

- Layoff: `Job loser/on layoff`, `other job loser`, `temporary job ended`
- Quit: `Job leaver`

2.3.2 Non-participation

Expanded questions on reasons nonparticipants left the labor force were added to the CPS in 1967 following recommendations in a 1962 report by the President's Committee to Appraise Employment and Unemployment. Subsequent research has argued that the answer to these questions is informative about future labor supply. For example, Deutermann Jr (1977) finds that nonparticipant prime age men who left their last job due to economic reasons or layoff are more likely to expect to return to the labor force within a year than those whose job ended for other reasons.[3]

The variable coding reason for leaving the last job is not easily available on CPS IPUMS for non-participants, those not actively searching for a job. This instead requires work with the raw CPS data. The next paragraphs will outline the process to distinguish separations into non-participation by reason of separations.

The question asked to individuals to inquire their reason of non-participation has slightly changed over the years, but is a close variant of:

`Why did ... leave that job?`

Before 1994, the question is asked to all non-participants who fulfill the following criteria: (1) currently not in the labor force, but worked for pay within the last five years, and (2) in the outgoing rotation group (ORG), which means the individuals are in month of sample 4 or 8. After 1994, the question is asked to individuals who (1) are currently not in the labor force, but worked for pay within the last 1 year, and (2) are in the outgoing rotation group (ORG). We restrict our sample to anyone who has worked in the past 12 months for the entire time period.[4] The possible answer choices to the question have changed over time and we harmonize the answers across all months and years and define a layoff or quit as follows:

- Layoff: `Temporary, seasonal or intermittent job completed`, `Slack work/business conditions`
- Quit: `Personal or family (including pregnancy)`, `Return to school`, `Health`, `Retirement or old age`, `Unsatisfactory work arrangements`

There are additional separations where the question asking the reason why the last job ended is not asked. These include, for example, retirements. We label these separations as other, but the reader should think of them as "unknown" since these separations such as retirements can certainly be preceded by an involuntary layoff as well as a planned quit.

[3] See Schwab (1974) for a men age 58-63.

[4] In theory, since we are looking at individuals who make a transition from employment in the previous month to non-employment in the current month, all individuals should fulfill this requirement, but a very small number reports not having worked in the past 1 year and we do not include them.

2.4 Linking over Time & Variable Construction

We employ linking and variable construction methodologies in order to come as close as possible to the construction used in IPUMS.

We follow Madrian and Lefgren (1999) when linking individuals across two consecutive months and verify match quality based on sex and age.[5] This method ensures that when we aggregate our flows to broad E-N and E-U rates, we recover the same transition probabilities that would be computed from IPUMS harmonized CPS data. In the CPS, the unique household and person identifier corresponds to the physical address of the individuals and therefore being able to match an individual does not necessarily imply matching the same individual but rather two individuals living at the same physical address in subsequent months. Personal characteristics, such as age and sex, which do not change over two subsequent months (or by not more than one in the case of age) and help to reduce false matches. Once we matched individuals across two subsequent months based on the above criteria, we use the matched data to compute the numbers of individuals in each labor market state in a given month.

For all labor market states with the exception of layoffs and quits into N, we simply count how many individuals are in each labor market state.[6] Since only individuals in the outgoing-rotation groups are asked about their reason for non-participation, we only have a subset of individuals responding to the question. We assume that the distribution of individuals by reason for non-participation is the same across all individuals in that month and use the share of quits and layoffs from the outgoing rotation groups multiplied by the total E-to-N transition rate to compute the number of layoffs for all other individuals making an employment to non-participation transition. Thus, we obtain flows numbers for individuals transitioning due to a layoff from E to N, and individuals transitioning due to a quit from E to N.

Once we have the numbers of individuals in each labor market state we compute flow rates between the different states. We compute the transition rates as the number of individuals with labor market state I in the previous month and labor market state J in the current month relative to all individuals with labor market state I in the previous month, such that

$$f_{IJ} = IJ_t / I_{t-1} \tag{5}$$

where $I = \{E, U, N\}$ and $J = \{E, U, N, UL, UQ, NL, NQ\}$ to obtain flow rates as shown in table 2.

Lastly, we seasonally adjust the data using the X-13ARIMA-SEATS seasonal adjustment program provided by the U.S. Census Bureau.

[5] We are not matching based on race since the answers have to this question has changed drastically over time.

[6] Consistent with best practices advise by IPUMS, we do not use weights in constructing the flow series.

2.5 Prime-age vs. Working-age population

All data results and statistics in the main text are for the prime-age population, those between 25 and 55 years old, only. We provide all results for the working-age population, those 16 years and older, in the appendix. We focus on the prime-age population because our main focus is on the labor supply decisions of marginal workers that are not necessarily driven by education or retirement choices.

3 Motivation

Figure 1 plots the full time series of the monthly quit and layoff rates to non-employment from 1978 - 2025 as a percent of total employment for the prime-age population in the United States. On average, 40% of all separations into non-employment are quits, and the remaining 60% are layoffs. This implies that in any given month, about 0.9% of all employed workers decide to leave their job. The majority of these quits are into non-participation, about 86% on average.[7]

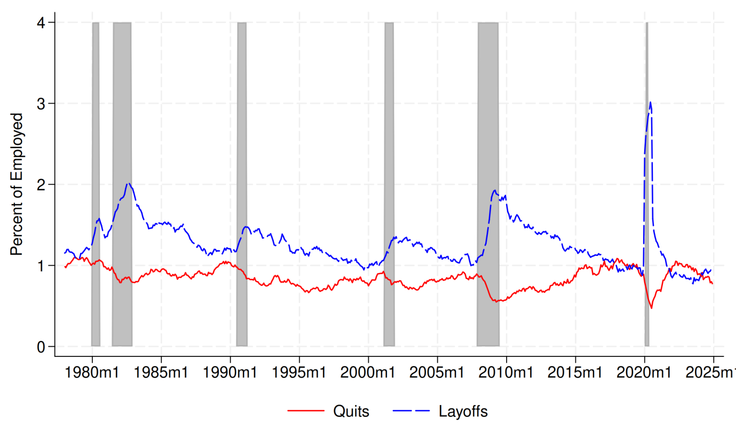

Figure 1: Monthly Quits and Layoff rates (as a percent of employment) (Monthly seasonally-adjusted data and 6-month centered moving average)

Since our data allows to classify all flows into non-participation as a layoff or quit, we can analyze the destination of all layoffs. Generally, the literature has assumed that an EU transition is due to a layoff, whereas an EN transition is due to a quit. We can see in figure 2 that 40% of all layoffs end in non-participation. Clearly, not all movements to non-participation are due to a quit decision as was previously assumed. This implies that about 40% of workers choose to leave the labor force after being laid-off.

[7]See appendix for figure.

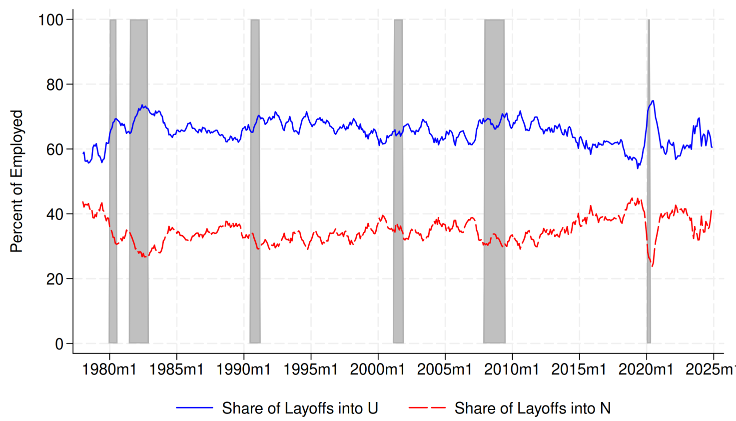

Figure 2: Share of layoffs by destination (Unemployment or non-participation)

In total, about 65% of all separations in a given month end in non-participation and are a result of labor supply decisions. Both types of labor supply decisions, quits and leaving the labor force after a layoff, tell us about the workers' value of employment compared to non-participation and help to inform labor supply elasticities.

The 40% of workers who quit in a given month are likely to have a negative employment surplus, or equivalently, the value of non-participation exceeds their value of employment. One interpretation of the large number of quits is that many individuals are on the margin of participation, i.e. the difference between their value of employment and non-participation is small. Additionally, we also observe large transition rates between non-participation and the labor force, about 10% of all individuals in non-participation join the labor force in the average month. These facts suggests that these workers are near the margin of participation. However, it is unlikely that all quits in the US are due to marginal workers. Some individuals who quit are likely to not join the labor force anytime soon and remain out of the labor force for a long time. The model will allow us to distinguish between the two types of quits. Understanding the distinction is important for estimating elasticities of labor supply. More importantly, the model will allow us to study the extensive labor supply decisions of these marginal workers, in particular, the factors that drive these decisions.

Similarly, 40% of all layoffs end in non-participation, which implies that these workers likely had a small but positive employment surplus, since they were working. However, the choice of non-participation after a layoff implies that these workers were very close to the margin. Our data, which allows for observing the reason and destination of a separation, is crucial to identify these marginal workers in a model. Laid-off workers and their decision to either stay in unemployment or leave the labor force is central to estimating labor supply elasticities in the model.

We can see in figure 2 that the share of layoffs into non-participation and unemployment display cyclical patterns. Table 3 shows how the destination of quits and layoffs vary over the business cycle. In times when the unemployment rate is high, the share of layoffs and quits into non-participation decline. The share of layoffs into non-participation is strongly negatively correlated with the unemployment rate. One interpretation of this finding is that laid off workers become more

Statistic	Share of Layoffs into N	Share of Quits into N
Corr(x, y)	-0.6222	-0.1253
SD(x)/SD(y)	0.0231	0.01543

Table 3: Business Cycle correlation of the share of layoffs and quits into N with the unemployment rate

attached to the labor force in recessions as more workers choose to remain unemployed after losing their job, for example, due to an increase in the value of unemployment benefits. Alternatively, this suggests that there is potential selection in the pool of laid-off workers in recessions. In normal times, when the share of laid-off workers exiting the labor force is larger, layoffs might be more targeted towards these small employment surplus marginal workers. In recessions, however, layoffs are less targeted and, therefore, the share of layoffs into non-participation declines. The model will help us to identify how much of the decline in the share of layoffs into non-participation is due to unemployment benefits versus a change in the composition of laid-off workers.

3.1 The Role of Composition Effects

We showed in the previous section that the share of individuals who leave the labor force after job loss declines during recessions. This section will investigate how much of this is due to a change in the composition of laid-off individuals in recessions or changes within different groups of laid-off workers. We do this by computing the counterfactual share of layoffs that flow into non-participation during each of the previous five recessions. Thus, we perform a shift-share analysis similar to Elsby et al. (2015) where share_{it}^{eln} is the weighted average of the fraction of laid-off workers exiting the labor force for different groups of laid-off workers:

$$\text{share}_t^{eln} = \sum_i \omega_{it} \text{share}_{it}^{eln} \tag{6}$$

where ω_{it} is the share group i among all laid-off workers in time t. In order to decompose the change in share_t^{eln}, we fix ω_{it} for each group i to the 12-month pre-recession average of ω_{it}, which we call ω_i^C. The counterfactual share of laid-off workers leaving the labor force, share_t^C, is computed as follows

$$\text{share}_t^C = \sum_i \omega_i^c \text{share}_{it}^{eln} \tag{7}$$

We calculate three different counterfactual analyses; by gender, race, and education, such that $i = \{\text{men}, \text{women}\}$, $i = \{\text{white}, \text{black}\}$, or $i = \{\text{high school at most}, \text{college and more}\}$,

Table 4 compares the actual change in the share of individuals who leave the labor force after job loss to the counterfactual change for each recession. We can see that in general the composition of the laid-off workers during recessions contributed little to the observed change in the share of laid-off leaving the labor force. Table 5 shows that while the importance of changes in the composition

Recession	Actual change in pp	Counterfactual change		
		Gender	Race	Education
Jan 1980 - Jul 1980	-0.0867	-0.0640	-0.0704	-0.0914
Jul 1981 - Nov 1982	-0.0870	-0.0657	-0.0926	-0.0921
Jul 1990 - Mar 1991	-0.0628	-0.0467	-0.0589	-0.0598
Mar 2001 - Nov 2001	-0.0508	-0.0502	-0.0556	-0.0507
Dec 2007 - Jun 2009	-0.1307	-0.1034	-0.1315	-0.1368

Table 4: Actual and counterfactual change (in pp) in the share of layoffs into non-participation

varies across the different groups and recessions, changes within groups contributes at least 75% to the decline in the share of layoffs into non-participation. Thus, the changes in labor supply decisions during recessionary periods is not a simple story of different people in the laid-off pool but rather a story about changes in people's labor supply decisions over the business cycle. It is worthwhile noting that our results differ from previous research[8] which has found that the pool of unemployed differs substantially in recessionary times from normal times. However, in our analysis we focus on everyone who is laid-off, i.e. we include individuals in unemployment and *non-participation*.

4 Model

We enrich a standard heterogeneous agent model of aggregate labor market dynamics to include features related to labor supply and selection. The foundation is an Aiyagari model, expanded to feature gross worker flows across three labor market states: employment (E), unemployment (U),

[8]See e.g. Elsby et al. (2015) and Mueller (2017)

Recession	Actual change in percentage	Counterfactual change		
		Gender	Education	Race
Jan 1980 - Jul 1980	-0.0867	-26.18%	-18.80%	+5.42%
Jul 1981 - Nov 1982	-0.0870	-24.48%	+6.44%	+5.86%
Jul 1990 - Mar 1991	-0.0628	-25.64%	-6.28%	-4.78%
Mar 2001 - Nov 2001	-0.0508	-1.18%	+9.45%	-0.19%
Dec 2007 - Jun 2009	-0.1307	-20.89%	+0.6%	+4.67%

Table 5: Actual and counterfactual change (in percentage) in the share of layoffs into non-participation

and nonparticipation (N) following the tradition of Krusell et al. (2017). To this we add a notion of quits and layoffs, each of which can happen at random or through endogenous choices of individuals thereby determining the extent of selection in job separation.

4.1 Environment

Time is discrete and the economy is populated by a unit measure of heterogeneous workers. The aggregate state is summarized by a vector $\mathcal{S}$ which includes a business cycle indicator (Z) and information necessary to determine equilibrium prices (w, r) and provide rational expectations over the future aggregate state.[9] Labor market frictions, including shocks related to job loss and job finding, as well as some government programs will depend on the business cycle indicator.

Preferences. Individuals have preferences represented by

$$E_t \sum_{t=0}^{\infty} \beta^t [\log(c_t) - \theta(s_t)], \tag{8}$$

where $c_t \geq 0$ is consumption in period t and $s_t \in \{0, 1\}$ is an indicators for whether the individual engages in active job search. The parameter $\theta > 0$ represents the disutility of active search and $0 < \beta < 1$ is the discount factor.

Budget Constraint and Earnings. Individuals cannot borrow but can accumulate assets a that offer a return $r(\mathcal{S})$. The budget constraint is:

$$c_t + a_{t+1} = (1 + r(\mathcal{S}))a_t + y_t$$
$$a_t \geq 0$$

Flow earnings, denoted by y_t, depend on the individual's labor market status. An employed worker's labor earnings are the product of her idiosyncratic labor productivity z_t and the market wage per efficiency unit of labor services $w(\mathcal{S})$, net of labor income tax $\tau(\mathcal{S})$:

$$y_t = (1 - \tau(\mathcal{S}))w(\mathcal{S}) \cdot z_t \tag{9}$$

As is typical, an individual's idiosyncratic productivity z_t is stochastic and follows an AR(1) process in logs:

$$\log z_{t+1} = \rho_z \log z_t + \varepsilon_{t+1}, \tag{10}$$

where the innovation ε_{t+1} is normally distributed with mean zero and standard deviation σ_ε.

Non-participants produce home production h, which is constant across individuals and time, and so their flow income is: $y_t = h$.

[9]TFP and the capital- effective labor ratio determine prices. Our computational approach follows Boppart et al. (2018) MIT-style shocks for expectations, as we detail later in the paper.

The model includes an Unemployment Insurance (UI) program that captures key features of the US system. To be eligible ($e = 1$) for UI, a worker must have previously been employed and experienced a layoff, as defined in the next section. Individuals who voluntarily quit their jobs are not eligible ($e = 0$). To receive benefits, an individual must engage in active search ($s = 1$) and incur the search cost θ. UI benefits have finite duration. An eligible individual loses their eligibility each period with probability $\mu(Z)$ which allows for UI extensions in recessions.

The value of the unemployment benefit is given by the following formula

$$b(z) = \begin{cases} b_0 z_t & \text{if } b_0 z_t \leq \bar{b} \\ \bar{b} & \text{otherwise} \end{cases}$$

where $\bar{b}$ is the maximum benefit amount and b_0 is the replacement rate.

Additionally, there is a means-tested consumption floor c_{min} representing safety net programs. This payment is available to all agents whose period resources $a_t + y_t$ fall below the consumption floor.

Production. The production function is Neoclassical. Output is produced by capital services K and effective labor L according to the function $A(Z)K^{\alpha_y}L^{1-\alpha_y}$. Total factor productivity, $A(Z)$, varies based on the aggregate state. The capital stock in production is equal to the total savings across households (i) $K = \int_i a_i$, and labor services are equal to the total efficiency units of labor supplied by the employed. Let $\mathcal{I}_e(i) = 1$ if agent i is employed, then this is $L = \int_i \mathcal{I}_e(i) z_i$. Capital and efficiency units of labor are rented at their respective marginal products: $r = A(Z)(\frac{L}{K})^{1-\alpha_y}$ and $w = A(Z)(\frac{K}{L})^{\alpha_y}$.

Government Budget. We assume that the government budget clears period by period. In any given period, total revenues equal the labor tax revenue from the employed $TR = \int_i \tau \mathcal{I}_e(i) z_i$, and total expenditures equal total spending on the consumption floor and unemployment benefits $TE = \int_i \mathcal{I}_{ub}(i) b_i + \int_i \mathcal{I}_c(i) c_{min}$ where $\mathcal{I}_{ub}(i) = 1$ for an unemployed individual eligible for unemployment insurance and $\mathcal{I}_c(i) c_{min} = 1$ for an individual receiving the consumption floor.

4.2 Labor Market Transitions and Employment Dynamics

Job Loss: Quits and Layoffs. Employed workers can lose their job each period in four different ways. Some we will classify as quits and others as layoffs. A portion of each of the quit and layoff hazards are totally exogenous and common to all workers. Another portion of each of the quit and layoff hazards are endogenous and selective: in equilibrium the hazard rates will be higher for workers with a lower surplus from employment. There is no on the job search and so any job loss results in a transition to non-employment.

First, the worker may quit. A selective quit is an endogenous quit that comes about when a worker assesses the value of non-participation or unemployment to be higher than continuing at work. Once separated he or she can choose whether to search for a new job in unemployment, or to move to non-participation for however long they like. There is also an exogenous shock p that forces a worker into non-participation until the shock goes away. We will call this type of quit "non-selective" or "random" and term the worker to be participation-constrained, meaning that they cannot make labor market choices and must stay in non-participation. This shock is meant to capture individuals very far from the margin of the labor market for reasons we don't model such as disability or the ample resources of a second earner. This shock follows a Markov process where $\pi^p(p', p)$ is the probability of moving to state p' from p. No worker exiting employment via a quit is eligible for UI.

Alternatively, the worker may be laid-off. As is standard in many labor market models, there is a probability $\delta(Z)$ that a worker is laid-off exogenously, and this probability is common across all workers but varies with the aggregate cyclical state (Z). We call these types of lay-offs "non-selective" or "random". The other type of lay-off is called "selective" and is generated by an endogenous response to a shock x. This shock has a random arrival rate of $\lambda_x(Z)$. Upon arrival, the worker must pay a utility cost x in order to remain employed. If they choose not to pay the cost, they are laid off. This is a simple way to capture selective firings of low-surplus workers within the firm. While both a selective lay-off and an selective quit target low-surplus workers, the layoff is completely transitory to the worker. It is meant to represent pressures idiosyncratic to the firm causing it to cut the least productive workers even though those workers could have positive surplus elsewhere, similar to a match quality shock. An endogenous quit, by contrast, is a separation initiated by a worker who does not have a positive employment surplus at *any* job. While all separated workers have the choice of whether to pay the search cost and enter unemployment or exit to non-participation, it is easy to show that since endogenous quits are due to a persistent decline in a worker's employment surplus, all selective quits will result in a transition to non-participation. By contrast, since endogenous selective layoffs are driven by transitory (one-period) changes in the employment surplus, they can result in the worker choosing either unemployment or employment.

The rich set of transitions between labor market states is key to our analysis, and so we summarize them here along with the clear ranking of average wage and duration of non-employment that will result from each type of flow in equilibrium. These rankings are validated by PSID analysis shown in Figure 24 (wages) and Figure 25 (duration) in the Appendix.

1. **Random Quits (*average wage, longest duration of non-employment*)**: With probability $\pi^p(1, 0)$, a worker is forced to quit into non-participation where they remain until the shock is removed. These are the quits of non-marginal workers who are far from participating

despite changes in policy or their circumstances.[10]

2. **Selective Quits (*lowest wage, medium-long duration of non-employment*)**: Employed individuals choose to quit if the surplus of employment relative to non-employment is negative. These are the quits of marginal workers who are relatively close to participating if policy or their circumstances change.

3. **Selective Layoffs (*medium-low wage, medium-short duration of non-employment*)**: With probability $\lambda_x(Z)$ a worker must pay a one-time utility cost x to continue an employment match. If they refuse, they are laid off. This will endogenously generate layoffs of marginal workers with low employment surplus who are more likely to choose to exit the labor force after the layoff.

4. **Random Layoffs (*average wage, shortest duration of non-employment*)**: With probability $\delta(Z)$, a worker is randomly laid off. This will equally generate layoffs of workers who are strongly attached, with high employment surplus, as it will generate layoffs of those marginally attached with low employment surplus.

Finally, a measure R_d of agents die each period and are replaced by an equal measure of newborn agents. The newborn agents begin non-employed and are endowed with assets equal to the average holdings in the economy. Deaths are not counted in labor market flows in the model (and they are not counted in the flows in the data). The newborn agents are counted as originating from non-participation. If they choose unemployment in their first period, then that will be counted as a flow from non-participation to unemployment.

Labor Market Frictions Both non-participants and the unemployed must wait to receive a job offer before moving to employment. The offer arrival rate for the unemployed is $\lambda_u(Z)$ and, without loss of generality, is greater than the arrival rate for the non-participants, $\lambda_n(Z)$. Both rates are assumed to be higher in expansions than in recessions.

4.3 Value Functions

The state variable of an individual is $(a, z, p; \mathcal{S})$, where a are assets, z is their idiosyncratic productivity, p is whether they are exogenously constrained from participation, and $\mathcal{S}$ is the aggregate state which includes the recession indicator Z. An unemployed worker will have an additional state for their UI eligibility, e.

[10] These quits are time invariant at the business cycle frequency and meant to capture non-participants who are highly inelastic at the business cycle frequency for reasons we do not model. In the quantitative model we provide these agents with flow consumption so they do not affect savings and this assumption becomes equivalent to if we had instead removed these flows from both model and data.

We express an individual's decision problem recursively using the following value functions. Let $V(\cdot)$ be the beginning of period value for an individual without an employment opportunity (either maintained employment from last period or has a job offer) and $J(\cdot)$ be the value for an individual with an employment opportunity; and $W(\cdot), U(\cdot), N(\cdot)$ be the values for an individual who has then chosen employed, unemployed, or a non-participant that period. For the employed:

$$W(a,z,p;\mathcal{S}) = \max_{c\geq 0, a'\geq 0} \log(c) + \beta E\Big[\underbrace{(1-\delta(Z)J(a',z',p',0;\mathcal{S}')}_{\text{No Layoff}} + \underbrace{\delta(Z)V(a',z'p',1;\mathcal{S}')}_{\text{Layoff}}\Big]$$

$$\text{st} \qquad c + a' = \max\{(1+r(\mathcal{S}))a + (1-\tau(\mathcal{S}))w(\mathcal{S})z, \underline{c}\}$$

For the unemployed:

$$U(a,z,p,e;\mathcal{S}) = \max_{c\geq 0, a'\geq 0} \log(c) - \theta + \beta E\Big[\underbrace{\lambda_u(Z)J(a',z',p',e';\mathcal{S}')}_{\text{Job Offer}} + \underbrace{(1-\lambda_u(Z))V(a',z',p',e';\mathcal{S}')}_{\text{No Job Offer}}\Big]$$

$$\text{st} \qquad c + a' = \max\{(1+r(\mathcal{S}))a + e\cdot b(z,\mathcal{S}), \underline{c}\}$$

$$e' = e \quad \text{with probability}(1-\mu(Z)); \qquad = 0 \quad \text{o/w}$$

For the non-participants:

$$N(a,z,p;\mathcal{S}) = \max_{c\geq 0, a'\geq 0} \log(c) + \beta E\Big[\underbrace{\lambda_n(Z)J(a',z',p',0;\mathcal{S}')}_{\text{Job Offer}} + \underbrace{(1-\lambda_n(Z))V(a',z',p',0;\mathcal{S}')}_{\text{No Job Offer}}\Big]$$

$$\text{st} \qquad c + a' = \max\{(1+r(\mathcal{S}))a + h, \underline{c}\}$$

The value for an individual without an employment opportunity depends on whether he or she is constrained to non-participation ($p = 0$) or not ($p = 1$). If not constrained, he or she can choose unemployment or non-participation freely each period.

$$V(a,z,p,e;\mathcal{S}) = \begin{cases} \max\{\underbrace{U(a,z,p,e;\mathcal{S})}_{\text{Unemployment}}, \underbrace{N(a,z,p;\mathcal{S})}_{\text{Non-participation}}\} & \text{if } p = 1 \\ N(a,z,p;\mathcal{S}) & \text{if } p = 0 \end{cases}$$

The value for an individual with an employment opportunity depends on whether or not he or she chooses to continue employment. The first line features the arrival of a match continuation shock x. Notice that a worker who receives the continuation cost shock but doesn't pay is laid-off and eligible for UI ($e = 1$), whereas the worker who quits is not eligible ($e = 0$).

$$J(a,z,p,e;\mathcal{S}) = \underbrace{\lambda_x(Z)[(1-p)\max\{\underbrace{W(a,z,p;\mathcal{S}) - x}_{\text{pay cost, keep employment}}, \underbrace{V(a,z,p,1;\mathcal{S})}_{\text{layoff}}\} + pN(a,z,p;\mathcal{S})]}_{\text{Match Continuation Cost Shock}}$$
$$+ (1-\lambda_x(Z))[(1-p)\max\{W(a,z,p;\mathcal{S}), \underbrace{V(a,z,p,0;\mathcal{S})}_{\text{quit}}\} + pN(a,z,p;\mathcal{S})]$$

4.4 Aggregate Shocks and Business Cycle Dynamics

Aggregate shocks to job finding and job loss and TFP are the primary candidate drivers of business cycles in the model. We model the cyclical shocks as a one-time, unanticipated aggregate impulse in the sense of Boppart et al. (2018). Let z_t be a scalar recession state that follows

$$z_t = \rho\, z_{t-1} + \eta_t, \qquad z_0 = 0, \qquad \eta_t = \mathbf{1}\{t = 1\}, \tag{11}$$

i.e. a single unit innovation hits at $t = 1$ and no further innovations are realized or anticipated, so $z_t = \rho^{t-1}$ for $t \geq 1$. Each cyclical parameter loads linearly on this common impulse, jumping to its bad-state value on impact and reverting geometrically to its good-state value,

$$\lambda_s^t = \lambda_s^* - \varepsilon_s\, z_t, \qquad s \in \{u, n\}, \tag{12}$$

$$\delta^t = \delta^* + \varepsilon_\delta\, z_t, \tag{13}$$

where λ_s^*, δ^* are the good-state (steady-state) values and $\varepsilon_s = \lambda_s^* - \lambda_s^B$, $\varepsilon_\delta = \delta^B - \delta^*$ are the impact responses, so that $\lambda_s^1 = \lambda_s^B$ and $\delta^1 = \delta^B$. The same construction applies to the continuation-cost arrival λ_x, the UI-extension rate μ, and aggregate TFP.

Notably, selective layoffs may also vary over the business cycle. The arrival of a job continuation cost shock x is $\lambda_x(Z)$, and may be higher or lower along the recession than in the steady state. We will see later that cyclicality of this parameter is only one determinant of whether the composition of layoffs moves towards higher or lower productivity workers in recessions.

Finally, to mimic the behavior of UI in the US in most business cycles since 1976, the probability of losing UI eligibility varies with the aggregate state and we assume $\mu^G > \mu^B$ to reflect the longer duration of benefits during recessions.

5 Computation

We solve the model with the impulse response methods of Boppart et al. (2018). The period length is one month. The steady state is the stationary equilibrium of the economy held at the normal-times (Z =G) parameter values: we solve the value functions and decision rules on a grid over assets and productivity, find the labor tax rate that clears the government budget, and compute the stationary distribution over states.

A recession is a single MIT vector impulse. Every cyclical parameter $(A, \delta, \lambda_u, \lambda_n, \lambda_x, \mu)$ jumps to its bad-state value in the impact month, and thereafter its deviation from the normal-times value decays geometrically at a common monthly rate ρ_x; the economy is then computed forward to convergence. Along this perfect-foresight path we guess (i) the path of the capital-effective labor ratio, which determines prices, and (ii) the path of the labor tax that clears the government budget; solve the household problem backward; simulate the distribution forward; and update the guesses until they converge. This technique is well-suited to our model because prices move little over the business cycle and the impact of tax rates on revenues is monotone within the range generated by all quantitatively relevant impulses we consider.

When we report a moment "at the recession trough" we mean its largest absolute deviation from steady state along the impulse. Since the model has no growth trend, we compare the data's peak-to-trough decline in real GDP per capita adjusted for trend growth over the same months.[11]

6 Estimation

The model period length is one month.

6.1 Targets and Protocol

We estimate the model by the Simulated Method of Moments in two stages. The nine steady-state parameters (θ, h, δ^G, λ_u^G, λ_n^G, λ_x^G, p_0, p_1, and x) are estimated to the normal-times averages (2002–2005) of nine level moments: the employment-population ratio, the unemployment rate, the EU, EN, UE, and NE transition rates, the quit rate, the layoff rate, and the share of layoffs to non-participation. These levels carry the cross-sectional message of the data: in particular, a 0.75% quit rate alongside a 3.5% flow back into the labor force can only be reconciled if a sizeable share of quits are of detached, non-returning workers; and a layoff-to-N share of 35.5% requires that layoffs disproportionately strike low-surplus workers.

The five recession parameters (A^B, δ^B, λ_u^B, λ_n^B, λ_x^B) are estimated to match the *business-cycle volatilities* of the Great-Financial-Crisis episode (2003–2019), rather than the levels at a single trough. The volatility targets are the standard deviations of the HP-filtered logs of six series: the employment-population ratio, the unemployment rate, quits, layoffs, the layoff share to non-participation, and labor productivity (output per worker).[12] UI does not expire during the recession ($\mu^B = 0$), matching the extensions to 99 weeks in the Great Recession; in normal times μ^G implies a mean benefit duration of six months.

The weighting matrix is judgmental. The employment-population ratio and the unemployment rate carry the heaviest weights, followed by the quit rate, layoff rate, and layoff share to N.[13] Since the empirical flows are not in steady state, a model that matched the exact EN, EU, UE, and NE

[11] Real GDP per capita fell 4.7% from 2007m12 to 2009m6; adding mean trend growth of 1.95% per year over those 18 months gives a 7.6% decline relative to trend, which we round to our 7.5% target. The same method gives 6.0% for the 1981–82 recession used in Section 11.

[12] The flow series are smoothed with a centered three-month moving average before filtering to remove month-to-month sampling noise.

[13] The weights on the level moments are 300 (E/pop), 100 (U rate), 5 (quits, layoffs, layoff share to N), and 3 on the rest, applied to squared proportional deviations of the normal-times levels.

flows would not match employment and unemployment and so we weight these moments less and prioritize getting the change in the macro aggregates correct. [14]

Several fixed parameters complete the calibration. The monthly interest rate delivers a 1% annualized return and $\beta = 1/(1+r)$. The labor share is $1-\alpha_y = 0.67$. The stationary variance of log productivity is 0.20.[15] The UI system provides a replacement rate of $b_0 = 0.5$ for earnings up to the median wage with the cap $\bar{b}$ binding thereafter. The consumption floor is $\underline{c} = 0.05$, approximately 5% of the median wage, capturing the Supplemental Nutrition Assistance Program, the primary welfare program utilized by temporarily unemployed workers (East and Simon (2024)).[16] The birth/death rate is $R_d = 0.5\%$ per month, which accounts for individuals leaving the survey without an observed separation and for new entrants who affect the NU, UE, and NE rates.

The Standard Model. Our straw man is the same economy stripped of the two ingredients we have introduced: selective layoffs ($\lambda_x = 0$, $x = 0$), so all layoffs are random; and random quits ($p_0 = 0$; $p_1 = 1$), so all quits are selective. This is closely aligned with workhorse models of labor market flows. The calibration also follows the standard approach in the literature by including the flows across states as targets (EU, UE, EN, NE) but excluding flows by reason (quits, layoffs, share of layoffs to non-participation).

6.2 Parameter Estimates

Table 6 reports the estimates. Three results deserve comment.

First, the two non-market flow values are moderate: home production is $h = 0.51$, about two-fifths of mean labor earnings, and the disutility of search is $\theta = 0.09$, equivalent to roughly a 9% consumption tax while searching. Together, they place a meaningful mass of low-productivity workers near the participation margin without putting anyone far above it, as is required to jointly match the levels of quits and of the layoff share to N.

Second, the full model splits separations into meaningful selective and random components. The continuation-cost shock arrives at rate $\lambda_x^G = 8.6\%$ per month and the cost $x = 1.04$ is roughly one month of median earnings; only workers whose surplus is below that threshold (about 4% of the employed) are released when it arrives, generating 31% of layoffs[17]. The participation shock removes $1 - p_1 = 0.25\%$ of unconstrained individuals per month and holds them out of the labor force for an average of $1/(1 - p_0) = 25$ months; it generates about a third of all quits. The random

[14] The computational appendix includes sensitivity analysis of the estimation, reports standard errors, and includes a full Jacobian analysis.

[15] This variance accounts for censoring due to selective non-participation, as discussed in Karahan and Ozkan (2013).

[16] East and Simon (2024) document that SNAP has the broadest eligibility among means-tested programs and shows the largest increase in receipt during unemployment, second only to UI.

[17] Workers are forward working and consider the total surplus, or present discounted values, of continuing work or not. It is not a comparison of current period flow values.

Parameter	Full model	Standard model
Estimated		
θ (disutility of active search)	0.095	0.073
h (home production level)	0.512	0.588
δ^G (random layoff, normal, %)	0.86	1.37
λ_u^G (offer rate, U, normal, %)	27.33	24.25
λ_n^G (offer rate, N, normal, %)	11.19	11.78
λ_x^G (cont.-cost shock, normal, %)	8.57	n/a
ρ_z (productivity persistence)	0.950	0.950
p_0 (remain participation-constrained)	0.960	0
p_1 (remain unconstrained)	0.998	1
x (continuation cost)	1.042	0 n/a
A^B (TFP in recession)	1.000	0.982
δ^B (random layoff, recession, %)	1.52	2.07
λ_u^B (offer rate, U, recession, %)	22.39	22.21
λ_n^B (offer rate, N, recession, %)	6.10	10.77
λ_x^B (cont.-cost shock, recession, %)	9.21	n/a
Fixed		
μ^G (UI expiry, normal, %)	10.91	10.91
μ^B (UI expiry, recession, %)	0.00	0.00
R_d (death/birth rate, %)	0.50	0.00
β (discount factor)	0.999	0.999
α_y (capital share)	0.330	0.330
b_0 (UI replacement rate)	0.500	0.500
$\bar{b}$ (UI cap)	1.000	1.000
$\underline{c}$ (consumption floor)	0.050	0.050
$\sigma^2_{\log z}$ (stationary variance)	0.200	0.200
Equilibrium Budget Balance		
τ^* (steady-state labor tax, %)	1.10	1.70

Table 6: **Model parameters** (monthly; probabilities in percent). Estimated parameters by SMM as described in the text. Standard errors and t-stats provided in the Appendix.

layoff rate $\delta^G = 0.86\%$ is well below the standard model's 1.37%: once selective layoffs exist, fewer random ones are needed.

Third, the recession is a *pure* labor-market event. It raises the random layoff rate by three-quarters (0.86% to 1.52%), raises the arrival of continuation-cost shocks modestly (8.6% to 9.2%), cuts the offer rate in unemployment by a sixth (27.3% to 22.4%), and cuts the offer rate in non-participation by nearly half (11.2% to 6.1%). The sharp fall in λ_n relative to λ_u matters for attachment: in recessions, leaving the labor force becomes a much worse route back to work than searching, which is part of why displaced workers stay. The TFP shock is essentially zero ($A^B \approx 1$). This is because layoffs in recessions target more highly productive workers than layoffs in the steady state causing effective labor to fall more than the number of employed. This channel is not operational in the standard model, which requires a decline in TFP ($A^B = 0.982$) to match GDP per worker volatility.

Moment	Data	Full model	Standard model
Levels: normal times (2002–2005)			
E/pop	79.4	78.7	78.3
Unemployment rate	3.40	3.39	3.38
EU rate	0.92	0.81	1.09
EN rate	1.48	1.16	1.44
UE rate	25.0	27.1	22.5
NE rate	7.0	6.4	7.2
Quit rate	**0.75**	**0.73**	**1.13**
Layoff rate	**1.20**	**1.24**	**1.40**
Share of layoffs to N	**35.5**	**34.6**	**22.1**
UN rate	18.0	10.9	10.2
NU rate	3.50	2.31	0.22
Business-cycle volatility: std of HP-filtered log, 2003–2019 (%)			
E/pop	0.91	1.03	0.91
Unemployment rate	11.0	14.3	11.0
Quit rate	**10.1**	**8.6**	**1.80**
Layoff rate	**11.0**	**8.4**	**6.7**
Share of layoffs to N	**8.6**	**8.0**	**4.46**
Labor productivity (GDP/worker)	0.71	0.50	0.71
Untargeted			
Mean short-N spell duration (months)	3.0	3.1	3.3

Table 7: **Model fit.** Moment sensitivity analysis provided in the Appendix.

6.3 The Standard Model Does not Produce the Composition of Separations or the Attachment Flows

Table 7 shows the fit of both models. The normal-times stocks and the targeted business-cycle volatilities of employment, the unemployment rate, and output per worker are targeted and matched by both, so the two share a normal-times labor market and a common macro volatility.[18] Both also target the four standard gross flows directly (EN, NE, EU, UE) and so the new moments surrounding quits and layoffs are where the models part ways.

The share of layoffs to N. In the standard model, 22.1% of laid-off workers exit the labor force compared to 35.5% in the data. All layoffs are random in the standard model which means they hit average workers, and average workers are more highly attached and more likely to enter unemployment and search. There is no parameter that can fix this: raising h to push more workers toward the margin would increase the quit rate, which is already too high. Matching the destination of layoffs *requires* that layoffs find the marginal workers, which is precisely what selective layoffs do. In the full model 31% of layoffs are selective, the laid-off marginal workers often exit, and the share of layoffs to N is 34.6%, much closer to the data. In the full model the mix layoffs shifts towards random layoffs that reach up the surplus distribution and allows the full model to replicate the volatility of exit after layoff. The standard model, by contrast, only generates half the volatility of exits after layoffs because the economic incentives encouraging counter-cyclical attachment are not strong enough (like those provided by UI extensions present in both the full and standard models).

Quits and layoffs. The standard model requires a higher quit rate than the data to match EN rates, 1.13% per month against 0.75%, because it produces too few layoffs to N. Quits also move the wrong way over the cycle. With no random component, every quit is an endogenous surplus quit; in a recession the value of holding a job falls, so *more* workers quit. This causes the quit rate to be counter-cyclical in the standard model, the opposite of the pro-cyclicality in the data. Quits are pro-cyclical in the full model because the hoarding response of marginal workers in response to falling job-finding rates drives quits down; but the standard model has no such margin since its marginal workers are already out of the labor force.

Attachment. The standard model has almost no flows from non-participation into search: its NU rate is 0.2% against 3.5% in the data, because every one of its non-participants is a worker who chose non-participation and sees no reason to search. While Krusell et al. (2017) and others solve this problem with exogenous idiosyncratic shocks to the cost of work, the full model we have developed generates two-thirds of these flows through fundamentals. The full model also reproduces

[18] The full model also provides a good fit to standard deviations re-cast as trough values such as a 75.7% recessionary trough Epop against 75.1% in the data.

the cyclicality of NU. During recessions, NU rises to 5.4% (data: 5.6%) as non-employed workers near the margin respond to the collapse of λ_n by choosing active search. UN also declines for this reason and through UI extensions. The standard model does not reproduce rising attachment: its NU rate inches up only from 0.2% to 0.5% (against 3.5% to 5.6% in the data).

6.4 What Quits, Layoffs, and Their Destinations Identify

This paper relies on quits, layoffs, and their destinations to identify the degree of selection of marginal workers in each type of employment separation. The inference rests on one piece of economics: *where* a separated worker goes reveals how much surplus the match had, and how that destination *moves* over the cycle reveals whether the separation was the worker's choice or the firm's.

Why the two margins matter. Labor supply and selection are the channels through which a given path of job finding and loss shocks generate different employment, unemployment, and welfare paths than mechanically implied. Labor supply (the participation and quit choices of workers near the margin) fixes the *sign* with which employment and unemployment respond to a shock. If marginal workers cling to jobs as job-finding collapses (hoarding) and the displaced keep searching (attachment), a recession destroys fewer matches than the shocks alone imply, and places more of the displaced in measured unemployment. In short, the recessionary response of labor supply buffers employment and amplifies unemployment. Selection (which workers are laid-off) affects how much output is destroyed per separation and where the displaced land. When layoffs target low-surplus workers each separation cleanses little output and sends more workers out of the labor force. When a recession shifts layoffs toward the randomly displaced, as in the full model, the same number of separations destroys more efficiency units and feeds disproportionally to unemployment. Both channels therefore move the depth of the employment decline, the meaning of an unemployment spike, the welfare cost of a recession, and the labor-supply distortion of unemployment insurance. The standard model omits selective layoffs entirely and produces a labor supply response of the wrong sign with rising quits in recessions which, as Section 8 shows, reverses the sign of how endogenous responses to the cycle contribute to aggregate volatility.

Levels identify the split of separations. Two normal-times moments do most of the work. The first is the *level* of quits together with the volume of re-entry from non-participation. A model in which every quit is a marginal labor-supply choice can match a quit rate of 0.75% per month only by placing a large mass of workers just below the participation margin, but those same workers churn back into search, so it cannot simultaneously match a low flow from non-participation into the labor force. The data show both a 0.75% quit rate *and* a 3.5% NU re-entry rate; reconciling them requires that a sizeable share of quits be of workers who are *not* near the margin and will not soon return: the random quits governed by (p_0, p_1). The quit level pins the total; the re-entry rate pins how much of it is the detached, non-returning kind. The second moment is the destination of layoffs, sN. A layoff that strikes a randomly chosen worker strikes an attached one, who searches, so a model of purely random layoffs sends almost everyone to unemployment ($sN \approx 17\%$). Matching the data's 35.5% requires that layoffs disproportionately find workers whose employment surplus is small enough that they prefer to leave the labor force. The level of sN therefore identifies the selective share of layoffs and, through it, the arrival rate λ_x and cost x of the continuation-cost shock: a higher sN demands that more layoffs be selective and that the surplus threshold x reach further into the employed distribution.

Cyclicality identifies the responses. How *big* labor supply and selection are over the cycle is identified by how the same moments move in recessions. The cyclicality of quits reveals the labor-supply response: holding decision rules fixed, a recession lowers the value of marginal jobs and would *raise* quits, so the observed *fall* of a third can only come from behavior (marginal workers hoarding jobs as outside options deteriorate), and the part of the quit decline that survives freezing decision rules (Section 8) measures the hoarding response directly. The cyclicality of sN reveals cyclical selection: part of its recession fall is increased labor supply (as marginal workers hold on, the layoffs that occur fall on slightly more attached workers), but the observed fall is larger than hoarding alone can produce, and the residual identifies a shift in the *incidence* of layoffs up the surplus distribution, the cyclical "upgrading" of the separating pool documented by Mueller (2017). Appendix E makes this quantitative with the moment Jacobian and a sensitivity matrix: λ_x and x load on sN, (p_0, p_1) load on the quit level and NU, and the recession frictions load on the cyclical movements of the same moments.

7 Sources of Business Cycles

With the estimated model in hand, we revisit the classic question of what drives recessions. Table 8 turns on one shock at a time (TFP, job finding (λ_u, λ_n), job loss (δ, λ_x), and the UI extension) and reports the trough deviation of each moment.[19] Agents understand and react to each counterfactual path.

Recessions are labor-market events. The GFC recession is carried by the labor-market shocks with essentially no role for TFP shocks in the full model. Contrary to Shimer (2012), the rise in unemployment in the full model is equally a job-loss and job finding event. This is partially because increases in job loss in the full model come with a shift from selecting low surplus, low

[19]Appendix Table 23 reports the same decomposition in volatility terms, as the standard deviation each shock generates on its own.

Moment	All shocks	TFP	Job finding	Job loss	UI extension
Full model					
Δ E/pop	-2.94	-0.01	-1.31	-1.56	-0.13
Δ U rate	+2.50	-0.03	+1.15	+1.24	+0.28
Δ EU rate	+0.63	-0.00	-0.03	+0.60	+0.05
Δ UE rate	-4.90	-0.00	-4.90	-0.00	-0.00
Δ EN rate	-0.49	+0.00	-0.55	+0.11	+0.01
Δ NE rate	-2.62	+0.01	-2.38	+0.16	-0.07
Δ UN rate	-5.93	-0.31	+4.51	-1.20	-3.32
Δ NU rate	+3.04	-0.08	+3.90	-0.18	-0.12
Δ Quit rate	-0.46	+0.00	-0.46	-0.01	-0.00
Δ Layoff rate	+0.60	+0.00	-0.10	+0.72	+0.05
Δ Share of layoffs to N	-12.97	+0.00	-4.63	-6.64	-1.30
GDP decline	2.56	0.03	1.18	1.13	0.10
Standard model					
Δ E/pop	-2.71	-0.78	-0.22	-1.49	+0.18
Δ U rate	+1.79	-0.19	+0.36	+0.91	+0.54
Δ EU rate	+0.47	-0.05	+0.06	+0.40	+0.02
Δ UE rate	-0.66	-0.00	-0.66	-0.00	-0.00
Δ EN rate	+0.36	+0.40	-0.36	+0.21	-0.04
Δ NE rate	-1.52	-0.07	-1.21	+0.12	-0.15
Δ UN rate	-6.73	+3.63	-2.98	+0.42	-5.95
Δ NU rate	+0.30	-0.14	+0.85	-0.03	-0.11
Δ Quit rate	+0.31	+0.34	-0.31	+0.09	+0.04
Δ Layoff rate	+0.52	-0.00	-0.00	+0.52	-0.00
Δ Share of layoffs to N	-3.77	+3.67	-4.16	+0.38	-1.72
GDP decline	5.94	5.29	0.38	1.36	0.04

Table 8: What drives recessions? Trough deviation of each moment from steady state (percentage points) along the recession impulse when only one shock is turned on at a time. "Job finding" moves (λ_u, λ_n); "Job loss" moves (δ, λ_x); "UI extension" sets $\mu^B = 0$.

productivity workers towards random layoffs of average workers. This generates a larger increase in unemployment per laid off worker. It also endogenously generates a decline in output-per-worker without requiring a fall in TFP. The standard model reads the same recession the other way around. It lacks a large increase in attachment in recessions, and so it leans on layoffs even more than a job-finding collapse to move unemployment. It lacks selective layoffs, and so it cannot generate output-per-worker volatility from composition and must assign a technology shock ($A^B = 0.98$) to carry the output-per-worker decline. Figure 3 plots the role of job-loss and job-finding as impulse responses. Other episodes are examined in Section 11.

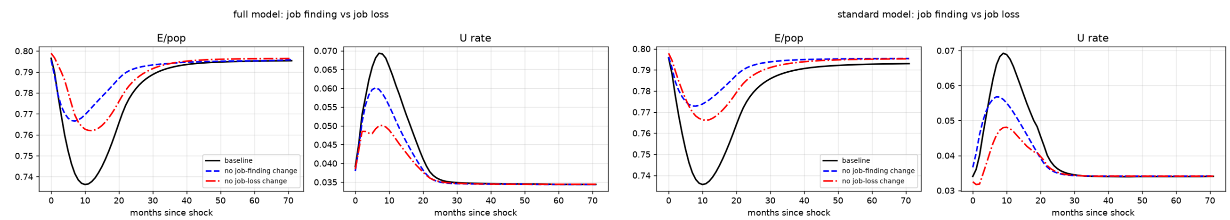

(a) Full Model: job finding vs job loss

(b) Standard Model: job finding vs job loss

Figure 3: **Role of job-finding and loss shocks in business cycle dynamics.** Baseline impulse response versus counterfactuals holding the indicated shock at its normal-times value.

Unemployment Insurance Extensions. UI extensions deliver 0.28 points (about a ninth of the total increase), which we examine in Section 9. The attachment flows show how policy and frictions interact. Job finding acting alone would push the UN rate up 4.9 points: with offers scarce and benefits expiring on schedule, unemployed workers near the margin give up and exit. The UI extension acting alone pushes UN *down* 2.9 points, holding the displaced in the labor force. In the full recession the extension wins: UN falls by 3.3 points, as in the data. Measured unemployment in a US recession is partially a product of policy-supported attachment in both models.

Quits are entirely job hoarding. The decline in quits during recessions in the full model can be entirely attributed to the decline in the job finding rate. When jobs are scarce, marginal workers quit less because they know it will be hard to find a job when their productivity drifts upwards in the future. The job hoarding response also exists in the standard model but is entirely offset by the decline in TFP which makes working less attractive and raises the incentive to quit.

8 Selection, Labor Supply, and the Marginal Workforce

This section quantifies the two channels the standard model lacks and identifies the workers through whom they operate.

8.1 How much selection is there, and when?

Share	Full model Normal	Full model Recession	Standard model Normal	Standard model Recession
Share of quits selective	66.2	19.0	100.0	100.0
Share of layoffs selective	30.6	17.7	0.00	0.00
Share of U stock entered selectively	3.28	2.01	0.15	0.39
Share of N stock entered selectively	44.9	42.2	73.5	74.8

Table 9: **Role of Selection: stocks and flows.** Selective quits are endogenous (negative-surplus) quits; selective layoffs are continuation-cost (x-shock) layoffs; the remainder are exogenous. A non-employed worker is counted as having entered selectively if the current non-employment spell began with a selective quit or layoff. Recession values are taken at the point of largest deviation during the episode.

Table 9 reports the selective share of each separation flow and of each non-employment stock. In normal times, 66% of quits and 31% of layoffs are selective, initiated by, or aimed at, low-surplus workers. Both shares fall sharply in recessions. The selective share of quits collapses to 19%, as marginal workers hoard their jobs causing the quits that remain to be dominated by the random quits of constrained workers. The selective share of layoffs falls to 18%, as the recession's separations are increasingly the random kind that hit attached workers. Whether selection should rise or fall in recessions is not obvious ex ante (the continuation-cost shock arrives *more* often in recessions, with λ_x rising from 8.6% to 9.2%), but fewer workers fail the continuation test because scarcity has raised the value of every job. The estimation needs exactly that mix to reproduce the joint fall in quits and in the layoff share to N. The model infers the composition of layoffs from labor supply behavior we observe in the data: since quits fall in recessions, the net surplus of marginal workers must be rising, and we should see *more* laid-off workers stay in the labor force; but the observed fall in the layoff share to N is larger than labor supply alone can deliver, so the incidence of layoffs must be shifting toward attached workers.

The standard model has no selective layoffs and no exogenous quits. None-the-less, some of the workers who quit endogenously make a later NU transition when their productivity drifts upwards but this is small. What is more common is for laid off workers to make a UN transition when they exhaust their UI, and this contributes to the roughly one-quarter of non-participants that did not enter via selective quit.

Table 10 tests the full model's selection patterns against the evidence in Mueller (2017), regressing composition measures on the unemployment rate along the impulse path. As in the data, separations shift strongly toward high-productivity workers in recessions: the productivity of workers transitioning from employment to unemployment rises with the aggregate unemployment rate with a coefficient of 2.6 in the full model which is close to Mueller's empirical estimate of 2.77. The standard model gives the opposite result with a negative coefficient. While Mueller did not

Dependent variable	Full model	Standard model	Mueller (2017)
Individual Productivity of Separated Workers			
z (unemployed)	2.60	-0.65	2.77
z (layoff)	2.09	0.20	
z (quit)	1.74	0.43	

Table 10: Cyclicality of composition: coefficients on the unemployment rate along the recession impulse path, in the spirit of Mueller (2017).

distinguish between quits and layoffs, we can see that both shift towards more highly productive workers as both become more random, but the change in the composition of layoffs drive the change in the composition of EU as workers who quit transition to N instead.

8.2 Who is marginal, and how many are there?

Marginal workers are those whose net employment surplus is near zero. They are the workers choosing selective quits and layoffs, which makes their labor supply choices the key elastic margin of the economy. The model identifies them sharply: a worker is at risk of a selective quit if her surplus is negative, and at risk of a selective layoff if her surplus is positive but below the continuation cost x. The likelihood of either falls with idiosyncratic productivity and rises with assets. For low-productivity individuals, home production and (if laid off) UI replace a large share of potential labor income, while high-asset individuals can use their savings to smooth consumption through non-employment. Two pieces of PSID micro-evidence, reported in Appendix C, support reading the destination of a separation as a signal of where a worker sits relative to this margin. Workers who quit to non-participation have the lowest pre-separation residual earnings and those laid off into unemployment the highest (Figure 24), exactly the surplus ordering the model assigns to selective quits versus the more attached workers caught by random layoffs; and non-employment spells are shortest for layoffs to unemployment and longest for flows to non-participation (Figure 25), consistent with destination-N separators being further from the participation margin. The finer flow classification, distinguishing destination as well as reason, therefore carries information about distance to the margin that a coarser quit/layoff split would discard.

Table 11 reports the distribution. In normal times the marginal workforce (negative surplus plus negative-if-cost) is 8.7% of the prime-age population, and another 5.2% are detached by the participation constraint. It is concentrated exactly where the data said it would be: 30.8% of non-participants are marginal and a further 28.0% constrained, against 3.7% of the employed and 2.7% of the unemployed. The recession *shrinks* the marginal workforce everywhere (among non-participants the negative-surplus group falls to 26.9%) as scarcity raises the value of holding (and of finding) a job and pulls workers back from the quit margin. By contrast, the share of the population with

Category	Total	Emp.	Unemp.	NILF
Full model				
Participation shock	5.2, 5.2	0.0, 0.0	0.0, 0.0	28.0, 26.3
Negative surplus	3.3, 2.7	0.0, 0.0	0.0, 0.0	17.6, 13.7
Negative surplus if continuation cost	5.4, 5.2	3.7, 3.4	2.7, 0.6	13.2, 13.2
Positive surplus always	86.1, 86.9	96.3, 96.6	98.3, 99.4	41.2, 46.8
Standard model				
Participation shock	0.0, 0.0	0.0, 0.0	0.0, 0.0	0.0, 0.0
Negative surplus	7.6, 7.9	0.0, 0.0	0.0, 0.0	39.8, 38.6
Negative surplus if continuation cost	0.0, 0.0	0.0, 0.0	0.0, 0.0	0.0, 0.0
Positive surplus always	92.4, 92.1	100.0, 100.0	100.0, 100.0	60.2, 61.4

Table 11: **Population by employment-surplus category (percent): normal times, recession trough.** The marginal workforce comprises the two middle (blue) categories: workers with negative surplus (they quit if employed) and workers whose surplus is positive but below the continuation cost x (they are released by a selective layoff if the shock arrives).

positive surplus actually falls in recessions in the standard model.

Figure 4 plots the surplus distribution over the unconstrained population in both models, in normal times and at the trough, with the quit region (< 0) and the selective-layoff region (0 to x) shaded. The full model places 9.1% of the unconstrained population in the two shaded regions, and in recessions the distribution shifts right and the regions thin out. The standard model has no selective-layoff region at all and, instead of shifting in a recession, the distribution of net surplus increases in variance.

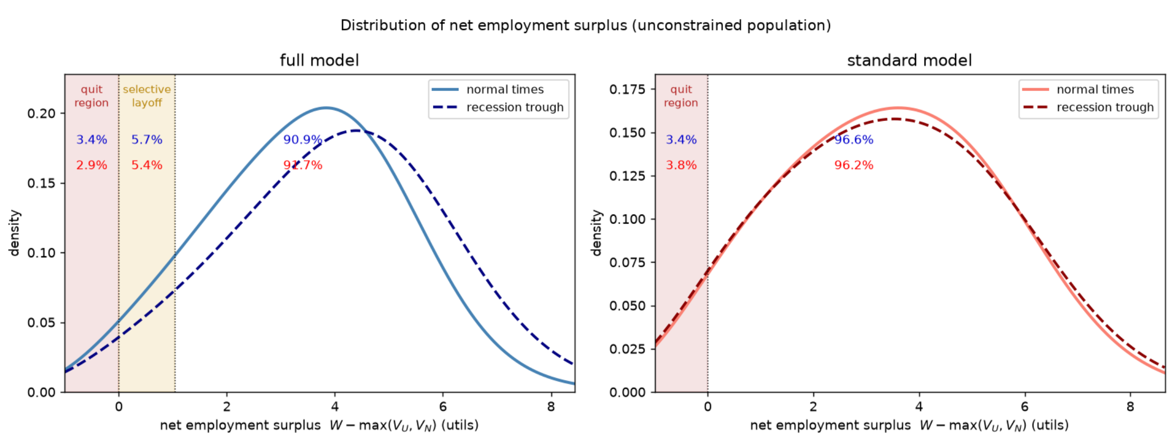

Figure 4: **Distribution of net employment surplus** ($W - \max(V_U, V_N)$, surplus relative to quitting) over the unconstrained population. Shaded: the quit region (< 0) and the selective-layoff region (0 to x). Region masses in blue (normal times) and red (recession trough).

8.3 What do labor supply and selection do to the cycle?

Table 12 reports the labor-supply and selection roles in steady-state to trough changes, and Appendix Table 24 reports them in volatility terms.

Moment	Full model				Standard model	
	Steady state	Baseline Δ	Labor supply role	Selection role	Baseline Δ	Labor supply role
E/pop	78.7	-2.94	+1.00	+0.72	-2.71	-0.44
U rate	3.39	+2.50	+0.32	+0.48	+1.79	+0.19
UN rate	10.9	-5.93	-2.24	-2.67	-6.73	-1.25
NU rate	2.31	+3.04	+3.14	+2.31	+0.30	+0.40
Quit rate	0.73	-0.46	-0.42	-0.30	+0.31	+0.35
Layoff rate	1.24	+0.60	+0.01	-0.14	+0.52	+0.00
Share of layoffs to N	34.6	-12.97	-3.23	-11.98	-3.77	-3.29
Want-to-work share	18.1	+3.53	-0.41	-0.42	+2.39	+0.12
GDP decline	–	2.56	-0.96	-0.10	5.94	+0.24

Table 12: **Roles of labor-supply response and selection in layoffs (percentage points).** Entries are trough deviations from steady state. "Labor supply role" is baseline minus a counterfactual holding decision rules and savings policies at steady state with the baseline incidence of selective layoffs fed in; a positive entry for E/pop means the labor-supply response BUFFERS the employment decline. "Selection role" is baseline minus a recession re-solved so the layoff share to N stays at steady state while matching the baseline layoff-rate increase.

Table 12 decomposes the recession. The labor-supply role compares the baseline to a counterfactual in which all decision rules that correspond to the steady state are used *except* the selective layoff choice, so the role isolates quit, search, and participation responses while keeping layoff incidence the same as the baseline. The selection role compares the baseline to a recession where the arrival rate of the continuation cost is re-solved so that the layoff share to N stays at its steady-state value while the total layoff rate path is unchanged.

Labor supply buffers the employment decline and amplifies unemployment. In the full model, the labor supply response makes employment fall *less*, by 1.0 percentage point at the trough, and makes the unemployment rate rise *more*, by 0.32 points. In other words, it buffers the employment decline by about a quarter and increases the unemployment rise by 13%. Two behaviors drive the counter-cyclical labor supply behind this result. First, job hoarding: essentially the entire 0.46 point decline in the quit rate is behavioral (the labor-supply role on quits is -0.42), keeping marginal workers employed who would have quit at steady-state policies. Second, added workers: the labor supply response more than accounts for the rise in the NU rate (a role of +3.14 against the +3.04 total rise), as non-employed workers respond to extended UI eligibility and the collapse of job finding rates in N. Labor supply accounts for only part of the higher attachment of laid-off workers via the decline in UN and smaller but still functioning share of layoffs to N. Figure

5 plots the labor supply role in the impulse response paths.

The standard model's labor supply has the opposite impact on employment. Its response *deepens* the employment decline (by 0.44 points) rather than cushioning it as its quits *rise* in the recession. With no selective layoffs and a productivity-driven recession that drives marginal workers' surplus negative, the only labor-supply margin it has is the discouragement of those marginal workers, who respond to the downturn by leaving. This is the sense in which the employment stabilizing impacts of labor supply choices are absent from the standard workhorse model.

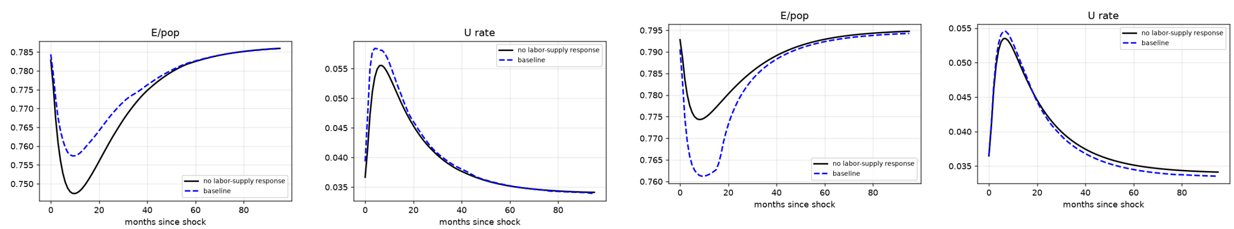

(a) E/pop, full model (c) U rate, full model (b) E/pop, standard model (d) U rate, standard model

Figure 5: **Role of labor supply in recessionary dynamics.** In each panel the dashed (baseline) line is the model's response and the solid line is the counterfactual with the labor-supply response shut off (decision rules frozen at steady state, baseline incidence of selective layoffs); the gap between them is the labor-supply role.

Selection works in the same direction, and is of similar magnitude. The selection role in Table 12 compares the baseline to a recession re-solved with the incidence of layoffs held at its normal-times composition. Cyclical selection buffers the employment decline by 0.72 points and amplifies the unemployment rise by 0.48, of the same order as the labor supply role. The mechanics: in the baseline, recessions shift layoffs toward attached workers, and attached workers *search* after displacement; in the selection-neutral counterfactual, the same share of layoffs keep falling on marginal workers, who exit. Selection also does most of the compositional work: without it the share of layoffs to N would fall only about a point in the recession. Figure 6 plots the selection-neutral impulse response.

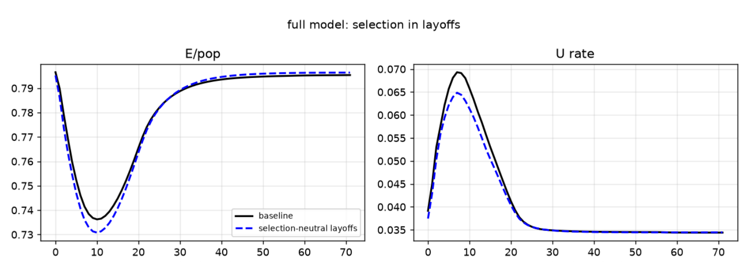

Figure 6: **Selection in layoffs.** Baseline impulse response versus the selection-neutral recession (layoff share to N held at steady state, layoff rate path matched).

Together, the two channels the standard model lacks (the labor supply of marginal workers and cyclical selection in layoffs) cushion the employment trough by 1.16 percentage points while adding 0.74 points to peak unemployment.[20] A reader who interprets a rise in unemployment one-for-one as a deterioration in employment will therefore over-read recessions through the lens of our model and under-read them through the lens of the standard one: in our economy, part of the unemployment spike *is* the stabilization working via workers staying attached and searching rather than leaving.

Why marginal workers stop quitting: "nowhere to go," not "wait for the pink slip" The recessionary decline in quits bundles two forces. A recession could make marginal workers hoard because the job finding has declined, or because rising job loss rates and longer UI benefits have made waiting for a layoff more attractive. We separate them by re-solving the recession with one force at a time switched off. Holding the job-finding path (λ_u, λ_n) at its steady-state level not only eliminates hoarding but *reverses* it: the labor-supply quit response flips from -0.28 to $+0.29$ points (quits now rise). Holding the recession UI extension (μ) at its steady-state level leaves the quit response essentially unchanged.

Dynamics of Marginal Workers over the Recession. To better understand the dynamics of labor supply over the cycle, we tag cohorts at the onset by their pre-shock employment surplus and follow how their employment change over the impulse response (Figure 7). In the left pane, we split the employed into the bottom and top quartiles of net surplus $W - \max(U_0, N)$. The marginal employed hoard (their quit hazard falls) causing their employment to increase briefly but overtime the increase in layoffs and decline in the job finding rate dominates and lowers their employment rate. The core employed never quit during the steady-state and have no off-setting job hoarding margin. They immediately have a decline in employment through higher job loss and lower job finding but eventually recover more quickly than the marginal employed because they are more likely to stay in the labor force and search more intensely for a job in unemployment.

9 The Welfare Cost of Business Cycles and the Value of UI

The previous sections established that (i) recessions displace workers less selectively, and (ii) marginal workers become more attached in recessions. This section asks how this changes our understanding of the welfare cost of recessions and the value of the UI extensions that accompany them. We compute, for every agent at the steady state, the permanent consumption-equivalent (CE) change that makes her indifferent between remaining at the steady state and entering the recession episode at date zero. Table 13 reports the results.

[20]These numbers are from an exercise where both channels are turned off simultaneously. Interaction causes the combined impact to be less than the sum of their individual impacts.

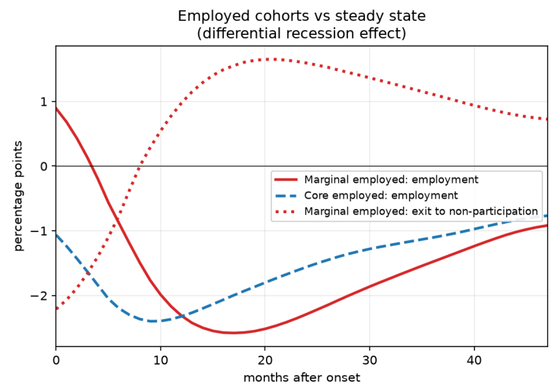

Figure 7: **Marginal-worker event study.** Change in employment rate of cohorts tagged at the recession onset and followed through the baseline impulse.

	Cost of recession		Value of UI extension	
	Full	Standard	Full	Standard
Mean	0.63	1.55	-0.120	-0.005
Std. dev.	0.24	0.15	0.052	0.045
Mean by labor market state at onset				
Employed	0.54	1.58	-0.126	-0.012
Unemployed	0.83	1.37	0.095	0.221
Non-participant	0.99	1.45	-0.123	-0.009
Cost with selection-neutral layoffs (full)	1.00	–	–	–

Table 13: **Welfare costs of recessions in consumption equivalents (%).**

The mean welfare cost of the recession episode in the full model is 0.63% of lifetime consumption. By labor market state at the onset, non-participants lose 0.99%, the unemployed 0.83%, and the employed 0.54%. That non-participants lose more than the employed is a direct consequence of the net value of employment increasing primarily through declines in the job finding rate. The employed lose more in the standard model due to the negative TFP shock that calibration requires.

Shifting selection in layoffs towards higher surplus workers actually makes recessions less costly. Re-running the recession with selection-neutral layoffs (the same aggregate layoff path falling randomly across workers rather than concentrating on those with the least to lose) raises the mean cost from 0.63% to 1.00%. This is because higher surplus workers are better self-insured: they have higher wages and more assets to better smooth consumption through job loss shocks.

The standard model overstates the cost of recessions. This is primarily because it requires a fall in TFP which hurts *every* worker in the economy whereas in the full model the only non-marginal workers who are affected are those who actually experience job loss, minus an increase in their UI tax rate. This is also why the the cost of recessions are flatter in the standard model. For distributional questions, and for any policy aimed at the displaced or the marginally attached, the models give different answers.

UI extensions. Our baseline recession includes the extension of UI to non-expiring benefits, as in the Great Recession. Solving the same recession with benefits expiring at their normal-times rate isolates the extension's value. The unemployed at the onset gain 0.12% of consumption: the insurance is real, and the extension is also what delivers the attachment dynamics we emphasized: it is the force that holds UN flows down and pulls non-participants into search. But the employed and non-participants each lose about 0.12%, paying through taxes and through the equilibrium of Section 7: the extension raises the unemployment rate by 0.3 percentage points and nudges the layoff rate up, as workers near the quit margin let themselves be laid off into insured unemployment instead. The population-weighted result is −0.12% in lifetime consumption: meaning that the average agent would pay to remove the extension.

The standard model scores the same policy as approximately neutral (−0.005%), because its marginal workers sit in non-participation and its employed are uniformly attached. There is hardly anyone whose unemployment-versus-non-participation choice the extension can move and employment barely responds to the policy. We temper this reading with the caveat that our experiment holds job creation fixed, omitting the vacancy response studied in Hagedorn et al. (2015) and the aggregate-demand effects studied elsewhere. What the model establishes is the size of the labor supply margin any such evaluation should consider.

10 Tracking Slack: Unemployed vs. Jobless who Want to Work

If unemployment conflates displacement, labor supply, and policy, what should we track instead? The model offers a natural answer: the jobless who *want to work*, that is non-employed individuals whose value of employment exceeds the value of their current state: $W > \max\{U, N\}$. This want-to-work (WTW) measure is the model-consistent stock of workers who would take a job today, cleansed of whether UI or search costs currently park them in U or in N.

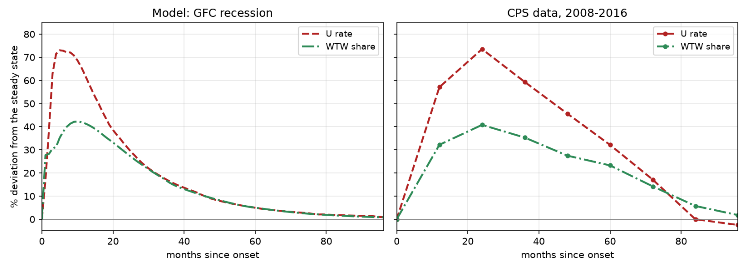

(a) Full Model

Figure 8: **Unemployment rate versus the want-to-work share over the recession.** Deviation as percent of the steady state along the impulse response. Left: Full Model. Right: CPS data 2008-2016.

Three facts stand out. First, WTW is larger than unemployment in the full model: 12% of the population is not employed, not participation constrained, and have a positive net surplus of employment (WTW), against an unemployment to population ratio of 2.9%. Clearly, unemployment does not capture all of the idle labor supply. Most of the difference is non-participants with positive employment surplus waiting on the slower offer rate λ_n rather than paying the cost of search. Second, WTW is a third as cyclical, shown in Figure 8 as steady-state (or peak) to trough deviations. The unemployment rate overstates the *proportional* deterioration in job-wanting and understates its breadth. Third, and most important for measurement: the wedge between the two series is partly labor supply response and partly changes in selection. Both the UI extension and less selectivity in layoffs lead laid-off workers to more often flow to U than N in recessions (as in the data), and an increase in NU (as in the data) draws workers with positive surplus into unemployment. Both changes in flows cause unemployment to rise more than the non-employed who WTW because the non-employed who WTW increasingly move to U not N.

The practical implication is that the unemployment rate is at its least informative exactly when it is watched most closely: in recessions with active UI policy. A rise in unemployment during an extension episode mixes displaced workers, marginal workers drawn into measured search by benefits, and laid-off workers who would have exited absent the extension. Our data series provide the raw material for tracking the alternative: the participation choices of laid-off workers (the ENL flow and the share of layoffs to N) reveal in close to real time where the margin of the labor market sits, and the model maps them into the stock of jobless who want to work. What we find in Figure 8 is that the CPS question asking non-participants if they want a job is consistent with that the model would predict given the flows, and so tracking responses to this question could be and additional useful indicator for real-time tracking of slack.

11 Episodes: 1980s, Great Financial Crisis, and 2020s

The baseline estimation disciplined the model with a single recession. The distribution of net employment surplus, however, is always changing overtime. For example, Ellieroth and Michaud (2024) argue that changes women's labor force attachment over the late 20th century fundamentally changed the response of job hoarding to the cycle. This means the labor supply backbone of our model is not time-invariant and needs to be re-calibrated to account for structural changes in participation elasticities.

We re-estimate the model on two further episodes with targets taken from CPS windows disciplining the steady-state to be contemporary to the episodes. First, we consider the 1980s recessions by re-estimating the full model with targets computed during the period 1976-1989. UI extensions are capped at a median duration of nine months (roughly 39 weeks), reflecting the less generous Federal Supplemental Compensation system of that era. For the 2020s, we consider the 2022-2025 period. The latter half of this episode is considered the steady-state and the roughly 2022-23 period provides a "booming" labor market. UI duration is held at 6 months over this period.

Table 14 reports the parameter estimates and fit for each episode. One important thing to notice is that it shows the range the model we developed can achieve. While TFP was estimated to be held constant in the GFC, this is not a tautology: the model requires a decline in TFP to generate the GDP-per-worker decline in the 1980s recessions. Table 15 decomposes each episode into the contribution of each shock and into the labor-supply and selection roles; Figure 9 plots the impulse responses.

The 1980s Recessions. The 1980s recession feature a larger role for selection and smaller role for labor supply. Job loss becomes a driver because we include with it the change in selection as job loss rates increase: it delivers 2.5 points of the decline in employment as the random layoff rate more than doubles (0.78% to 1.93%). While the drop in TFP is 3.1 ppt, its decline contributes only 0.65 points to the employment decline. The era's less generous UI also changes what the unemployment rate measures as the impact of UI policy becomes negligible. Instead the 1980's rise in unemployment was overwhelmingly displaced workers searching, not benefit-supported attachment, yet attachment still rose, as in the data: the model reproduces the UN decline (13.2% to 11.2%; data 19.6% to 15.4%) and the NU rise (3.0% to 3.9%; data 2.8% to 3.9%) of that episode, driven there by job scarcity

	GFC (baseline)	1980s	2022–23 boom
Parameters: steady state, episode state			
θ (disutility of active search)	0.095	0.093	0.090
h (home production level)	0.512	0.504	0.509
δ^G (random layoff, %)	0.86	0.78	0.85
λ_u^G (offer rate, U, %)	27.33	28.38	29.08
λ_n^G (offer rate, N, %)	11.19	10.49	10.63
λ_x^G (cont.-cost shock, %)	8.57	8.41	7.61
x (continuation cost)	1.042	1.087	1.082
A^B (TFP)	1.000	0.969	1.007
δ^B (random layoff, %)	1.52	1.93	0.76
λ_u^B (offer rate, U, %)	22.39	21.63	29.08
λ_n^B (offer rate, N, %)	6.10	5.36	13.26
λ_x^B (cont.-cost shock, %)	9.21	9.97	5.33
UI in episode	no expiry	39 weeks	26 weeks
Model outcomes: normal $\rightarrow$ episode			
E/pop	78.7 $\rightarrow$ 75.7	77.2 $\rightarrow$ 72.6	80.0 $\rightarrow$ 81.0
U rate	3.39 $\rightarrow$ 5.9	3.51 $\rightarrow$ 6.5	3.44 $\rightarrow$ 3.03
Share of layoffs to N	34.6 $\rightarrow$ 21.6	36.8 $\rightarrow$ 22.4	33.5 $\rightarrow$ 31.1
Share of layoffs selective	30.6 $\rightarrow$ 17.7	35.2 $\rightarrow$ 17.9	30.7 $\rightarrow$ 26.2
UN rate	10.9 $\rightarrow$ 5.0	13.2 $\rightarrow$ 11.2	11.7 $\rightarrow$ 15.0
NU rate	2.31 $\rightarrow$ 5.4	2.96 $\rightarrow$ 3.89	3.08 $\rightarrow$ 2.15
Want-to-work share	18.1 $\rightarrow$ 21.6	20.1 $\rightarrow$ 24.8	17.2 $\rightarrow$ 16.0

Table 14: **Episodes: calibration and fit**.

	GFC (baseline)	1980s	2022–23 boom
Contribution to episode Δ U rate (acting alone, pp)			
Job loss	+1.24	+2.04	+0.25
Job finding	+1.15	+1.25	-0.61
TFP	-0.03	+0.04	+0.01
Contribution to episode Δ E/pop (acting alone, pp)			
Job loss	-1.56	-2.54	+0.70
Job finding	-1.31	-1.04	+0.16
TFP	-0.01	-0.65	+0.11
Impact of labor supply and selection (pp at trough)			
Labor supply, E/pop	+1.00	+0.20	-0.29
Selection, E/pop	+0.72	+1.05	+0.27
Labor supply, U rate	+0.50	+0.21	-0.16
Selection, U rate	+0.55	+0.78	-0.17

Table 15: **Episodes: recession decomposition.** The top panels give each shock's contribution, acting alone, to the episode change in the unemployment rate and Epop. The bottom panel gives the impact of the labor-supply response and of cyclical selection at the trough (boom peak for the 2020s), each measured as the baseline outcome minus the counterfactual with that channel shut off; a positive E/pop entry means the channel holds employment up.

rather than policy.

The 1980s recession is more impacted by the change in selection in part because the 1970's-1980's steady state features more selection to start with: 37% of normal-times layoffs flow to non-participation (the CPS target is 41%), and the recession undoes it: the selective share of layoffs collapses as random layoffs take over, and the share of layoffs to N falls to 22% (target 29%). The cyclical "upgrading" of the separating pool documented by Mueller (2017) is at full force in the model's 1980s.

The 2020s Expansion. The model interprets the 2020s expansion as held back by a significant decline in labor supply. For context, the decline in labor supply held back Epop growth by 0.3 ppt which is about a third of the actual increase during these years. Figure 9 shows that the decline in unemployment considerably understates the decline in idle labor supply: the share of non-employed who wanted a job fell by much more. This is exactly the divergence between slack and measured unemployment that Section 10 emphasizes, now running in the expansionary direction.[21]

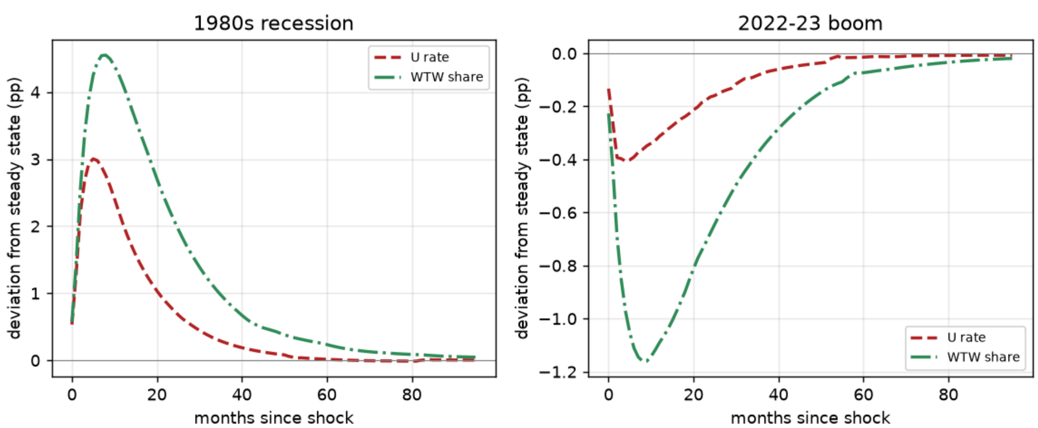

Figure 9: **Episodes.** Impulse responses of the unemployment (U) rate and the non-employed who want-to-work (WTW) share as deviations from each episode's own steady state: 1980s recession, and 2022–23 boom.

Figures 10 and 11 visualize the labor-supply and selection decomposition into each episode. In every episode the dashed (no labor supply) counterfactuals sit outside the baseline for employment and inside it for unemployment, implying qualitatively the same stabilizing pattern as in the GFC. Yet each episode is distinct as labor supply plays a much smaller role in the 1980s recession. The dash-dotted (no selection) line follows the same stabilizing pattern for the GFC and 1980s, but not for employment in the 2022-23 boom. Selection moved towards more random layoffs in 2022-23, a pattern we typically saw in recessions. This may suggest that the relationship between labor market

[21] The job-to-job quit surge of the Great Resignation is a separate phenomenon: in our model the cyclical force on quits is quits of marginal workers to non-employment, whereas the 2021–22 quits researchers studied were without non-employment spells (Bagga et al. (2023), Cai and Heathcote (2023), Şahin and Tasci (2022), Michaels (2024)).

tightness and selectivity in layoffs is non-monotone, or that the 2020's were a special episode for a variety of reasons.

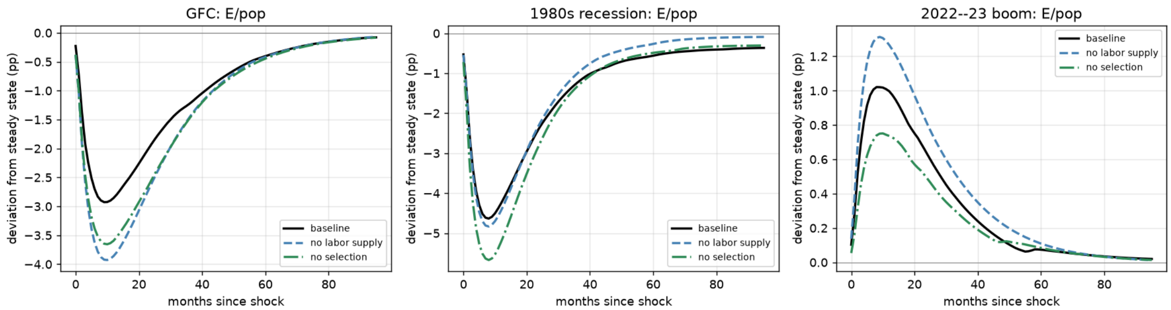

Figure 10: **Labor supply and selection in employment, by episode.** E/pop response of each episode (deviation from its own steady state): the baseline, the counterfactual with the labor-supply response shut off, and the counterfactual with selection-neutral layoffs. The gap from the baseline to each counterfactual is that channel's contribution.

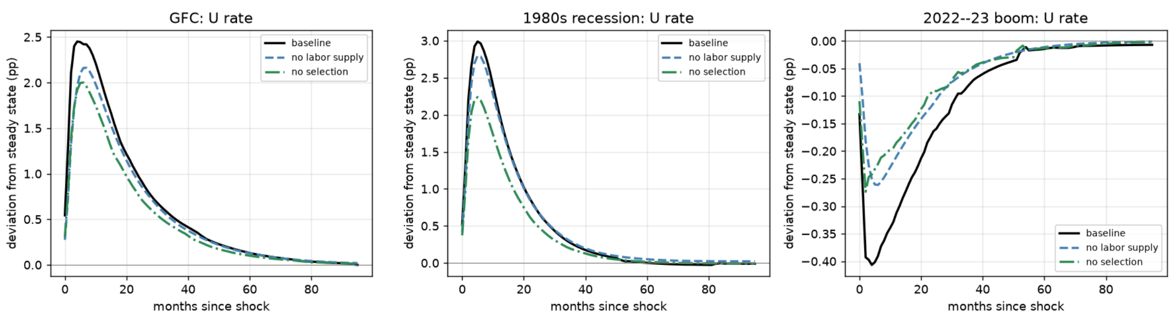

Figure 11: **Labor supply and selection in unemployment, by episode.** As in Figure 10, for the unemployment rate.

The episodes also sharpen the measurement lesson of Section 10. In the late 1970s the marginal workforce was bigger, yet unemployment and WTW track each other more closely over the 1980s cycle because the modest UI policy introduced little wedge between searching and wanting to work. In the 2022–23 boom they diverge for the opposite reason: the unemployment rate falls 0.4 points while the WTW share falls a full point and employment rises 1.1: the tight market absorbed jobless workers, most from outside measured unemployment, exactly the margin the unemployment rate misses. Across episodes, the unemployment rate is a good proxy for slack precisely when UI policy is quiet and the marginal workforce is small; when either condition fails, the flows we construct, in particular the destination of laid-off workers, are the observable correction.

12 Conclusion

The composition of separations is a demanding new test for business-cycle models of the labor market, and the workhorse model fails it. Confronted with monthly CPS series on quits and layoffs by destination, a model with random layoffs and purely endogenous quits produces too many quits, sends laid-off workers to unemployment when a third of them walk out of the labor force, and, because its only labor-supply margin is discouragement, gets backwards the direction in which attachment moves over the cycle. Two ingredients repair it: layoffs that select on employment surplus, and quits that strike workers far from the participation margin. The data identify both without strain: the level of quits together with the volume of re-entry from non-participation pin down the random quits, and the destination of laid-off workers pins down the selection in layoffs.

Getting the composition right overturns what labor supply does over the cycle: a recession is partly stabilized from within. Recessions make the jobless more attached—quits fall, exits from unemployment to the sidelines fall, and entry from the sidelines into search rises—and this works through two channels the standard model misses or signs incorrectly. Labor supply buffers the employment decline by 1.0 percentage point at the trough while adding 0.32 points to measured unemployment, as marginal workers hoard their jobs and the displaced keep searching. Cyclical de-selection of layoffs works in the same direction and is of similar magnitude, cushioning employment by a further 0.72 points as layoffs shift from the marginal toward the firmly attached. Together the two channels lift the employment trough by 1.16 percentage points and raise peak unemployment by 0.74. Part of the unemployment spike is therefore the stabilization working—workers staying in the game rather than leaving it—and a reader who equates the rise in unemployment one-for-one with lost employment misreads the recession.

These mechanics reorganize what drives the cycle and what it costs. The employment downturn is a labor-market event, carried by job finding and job loss rather than technology; the output cycle is its mirror image, driven by the changing composition of the employed, since hoarding spares low-productivity workers and output per worker falls without a technology shock. The welfare cost is concentrated on the non-employed: a recession costs 0.63% of lifetime consumption on average, but non-participants and the unemployed at the onset (losing 0.99% and 0.83%) bear far more than the employed (0.54%). This is a lower cost than a standard calibration that relies on a large TFP decline which adds an additional cost to everyone. The unemployment-insurance extensions that accompany modern recessions primarily finance attachment rather than distort search: they hold the displaced in measured unemployment and pull non-participants into it, so that the average worker would pay 0.12% of consumption to remove them even though the unemployed value the insurance.

Finally, these same forces make the unemployment rate a noisy gauge of slack precisely when it is watched most closely, in recessions with active UI policy, because it conflates displacement, labor supply, and benefits. The model points to a cleaner object that the data construct in close to real time—the jobless who want to work—and the episodes show why it matters: the displacement-heavy recessions of the early 1980s, where attachment rose on job scarcity rather than policy, and the tight labor market of 2022–23, where the unemployment rate fell by a fraction of the decline in idle labor supply. We are committed to updating our monthly CPS quit and layoff series so that researchers and policymakers can use the destination of separating workers, and the framework that prices it, to look beyond standard aggregates and track the margin of the labor market as it moves.

References

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

Ahn, Hie Joo and James D Hamilton, "Heterogeneity and unemployment dynamics," *Journal of Business & Economic Statistics*, 2021, *39* (3), 798–811.

_ , Bart Hobijn, and Ayşegül Şahin, "The dual US labor market uncovered," Technical Report, National Bureau of Economic Research 2023.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro, "Measuring the sensitivity of parameter estimates to estimation moments," *The Quarterly Journal of Economics*, 2017, *132* (4), 1553–1592.

Bagga, Sadhika, Lukas Mann, Aysegül Sahin, and Giovanni L Violante, "Job Amenity Shocks and Labor Reallocation," 2023.

Barnichon, Regis and Andrew Figura, "Labor market heterogeneity and the aggregate matching function," *American Economic Journal: Macroeconomics*, 2015, *7* (4), 222–49.

Bednarzik, Robert W and Deborah P Klein, "Labor force trends: a synthesis and analysis," *Monthly Labor Review*, 1977, *100* (10), 3–12.

Blanchard, Oliver Jean and Peter Diamond, "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity*, 1990, *2*, 85–143.

Boppart, Timo, Per Krusell, and Kurt Mitman, "Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative," *Journal of Economic Dynamics and Control*, 2018, *89*, 68–92. Fed St. Louis-JEDC-SCG-SNB-UniBern Conference, titled: "Fiscal and Monetary Policies".

Cai, Zhifeng and Jonathan Heathcote, *The great resignation and optimal unemployment insurance*, Federal Reserve Bank of Minneapolis, Research Division, 2023.

Cairó, Isabel, Shigeru Fujita, and Camilo Morales-Jiménez, "The cyclicality of labor force participation flows: The role of labor supply elasticities and wage rigidity," *Review of Economic Dynamics*, 2022, *43*, 197–216.

Cheng, Edmond, Nicole Hudson, and Jurgen Kropf, "The CES/JOLTS Divergence: How to Apply the Monthly Alignment Method to Help Close the Gap," Technical Report, U.S. Bureau of Labor Statistics 2009.

Chodorow-Reich, Gabriel and Loukas Karabarbounis, "The cyclicality of the opportunity cost of employment," *Journal of Political Economy*, 2016, *124* (6), 1563–1618.
ΩŞahin and Tasci

Şahin, Ayşegul and Murat Tasci, "The Great Resignation and the Paycheck Protection Program," *Economic Commentary*, 2022.

East, Chloe N and David Simon, "The safety net and job loss: How much insurance do public programs provide?," *Journal of Public Economics*, 2024, *238*, 105171.

Ellieroth, Kathrin and Amanda Michaud, "From Trend to Cycle: the Changing Careers of Married women and Business Cycle Risk," *Working Paper*, 2024.

Elsby, M., B. Hobijn, and Ayşegul Şahin, "On the Importance of the Participation Margin for Labor Market Fluctuations," *Journal of Monetary Economics*, 2015, *72*, 64–82.

Elsby, Michael, Bart Hobijn, Fatih Karahan, Gizem Koşar, and Ayşegul Şahin, "Flow Origins of Labor Force Participation Fluctuations," *American Economic Reviews Papers and Proceedings*, 2019.

Faberman, R Jason, "Analyzing the JOLTS hires and separations data," in "Proceedings of the" Citeseer 2005.

Farber, Henry S and Robert G Valletta, "Do extended unemployment benefits lengthen unemployment spells? Evidence from recent cycles in the US labor market," *Journal of Human Resources*, 2015, *50* (4), 873–909.

Flaim, Paul O, "Persons not in the labor force: who they are and why they don't work," *Monthly Labor Review*, 1969, pp. 3–14.

_ , "Discouraged workers and changes in unemployment," *Monthly Labor Review*, 1973, pp. 8–16.

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

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, Grace Cooper, Stephanie Richards, Megan Schouweiler, and Michael Westberry, "IPUMS CPS: Version 11.0 [dataset]," Technical Report, Minneapolis, MN: IPUMS 2023. https://doi.org/10.18128/D030.V11.0.

Fujita, Shigeru, Giuseppe Moscarini, and Fabien Postel-Vinay, "Measuring Employer-to-Employer Reallocatio," *American Economic Journal: Macroeconomics*, Forthcoming.

Garibaldi, Pietro and Etienne Wasmer, "Equilibrium Search Unemployment, Endogenous Participation, and Labor Market Flows," *Journal of the European Economic Association*, 2005, *3*, 851–882.

Gellner, Christopher G, "Enlarging the concept of a labor reserve," *Monthly Lab. Rev.*, 1975, *98*, 20.

Gibbons, Robert and Lawrence F Katz, "Layoffs and lemons," *Journal of Labor Economics*, 1991, *9* (4), 351–380.

Graves, Sebastian, Christopher K Huckfeldt, and Eric T Swanson, Technical Report, National Bureau of Economic Research 2023.

Gregory, Victoria, Guido Menzio, and David Wiczer, "The alpha beta gamma of the labor market," *Journal of Monetary Economics*, 2025, *150*, 103695.

Hagedorn, Marcus and Iourii Manovskii, "The cyclical behavior of equilibrium unemployment and vacancies revisited," *American Economic Review*, 2008, *98* (4), 1692–1706.

_ , _ , and Kurt Mitman, "The impact of unemployment benefit extensions on employment: the 2014 employment miracle?," Technical Report, National Bureau of Economic Research 2015.

Hornstein, Andreas, Marianna Kudlyak, and Fabian Lange, "Unemployment and vacancy dynamics with imperfect financial markets," *Economic Quarterly*, 2014, *100* (2), 135–153.

Jones, Stephen RG and W Craig Riddell, "The measurement of unemployment: An empirical approach," *Econometrica*, 1999, *67* (1), 147–162.

_ and _ , "Unemployment and non-employment: Heterogeneities in labor market states," *Review of Economics and Statistics*, 2006, *88* (2), 314–323.

Jr, William V Deutermann, "Another look at working-age men who are not in the labor force," *Monthly Lab. Rev.*, 1977, *100*, 9.

Karahan, Fatih and Serdar Ozkan, "On the persistence of income shocks over the life cycle: Evidence, theory, and implications," *Review of Economic Dynamics*, 2013, *16* (3), 452–476.

Koşar, Gizem and Wilbert Van der Klaauw, "Workers' perceptions of earnings growth and employment risk," *FRB of New York Staff Report*, 2023, (1056).

Krueger, Alan B, Judd Cramer, and David Cho, "Are the long-term unemployed on the margins of the labor market?," *Brookings Papers on Economic Activity*, 2014, *2014* (1), 229–299.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegul Şahin, "Gross Workers Flows over the Business Cycle," *American Economic Review*, 2017, *107(11)*, 3447–3476.

Ljungqvist, Lars and Thomas J Sargent, "The fundamental surplus," *American Economic Review*, 2017, *107* (9), 2630–2665.

Madrian, Brigitte C. and Lars John Lefgren, "An Approach to Longitudinally Matching the Current Population Survey," *NBER Technical Working Paper No. 247*, 1999.

Michaels, Ryan, "What Explains the Great Resignation?," Technical Report, Federal Reserve Bank of Philadelphia 2024.

Mitman, Kurt and Stanislav Rabinovich, "Do Unemployment Benefit Extensions Explain the Emergence of Jobless Recoveries?," 2019.

Morchio, Iacopo, "Work histories and lifetime unemployment," *International Economic Review*, 2020, *61* (1), 321–350.

Mueller, A., "Separations, Sorting, and Cyclical Unemployment," *American Economic Review*, 2017, *107* (7), 2081–2107.

Mui, Preston and Benjamin Schoefer, "Reservation raises: The aggregate labor supply curve at the extensive margin," *NBER Working Paper*, 2021, (w28770).

Mukoyama, Toshihiko, Christina Patterson, and Aysegul Sahin, "The cyclical behavior of unemployment and wages under information frictions," *Journal of the European Economic Association*, 2018, *16* (3), 825–856.

Qiu, Xincheng, "Vacant jobs," *University of Pennsylvania*, 2022.

Rothstein, Jesse, "Unemployment insurance and job search in the Great Recession," *Brookings Papers on Economic Activity*, 2011, *2011* (2), 143–213.

Schwab, Karen, "Early labor-force withdrawal of men: participants and nonparticipants aged 58-63," *Soc. Sec. Bull.*, 1974, *37*, 24.

Shimer, Robert, "The cyclical behavior of equilibrium unemployment and vacancies," *American economic review*, 2005, *95* (1), 25–49.

_ , "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 2012, *15*, 127–147.

Simmons, Michael, "Job-to-job transitions, job finding and the ins of unemployment," *Labour Economics*, 2023, *80*, 102304.

Tauchen, George, "Finite state Markov-chain approximations to univariate and vector autoregressions," *Economics Letters*, 1986, *20* (2), 177–181.

A Data Robustness

A.1 Choice of moving average filter

In the main text, all timeseries were smoothed using a 6-month centered moving average smoother. We will show in the following that the choice of the smoothing parameter as well as whether it is centered or not does not change the data in any significant way. The following figures plot our layoff series, quit series, and total separations for four different smoothing techniques. 6-month centered is the standard we use in the main text, which means we include the previous 3 months, the current month, and three forward terms. 3-month, 4-month, and 6-month only include the previous 3, 4, and 6 months respectively, as well as the current month.

We see that the different lengths really only affects the pandemic period as it was so short but so extreme. It does not seem to affect other recessions or expansions. We checked including both lags and leads versus only including leads to make sure the most recent data is not significantly affected by the moving average filter. As we can see in the following figures, we see no difference between the two methods for the most recent observations.

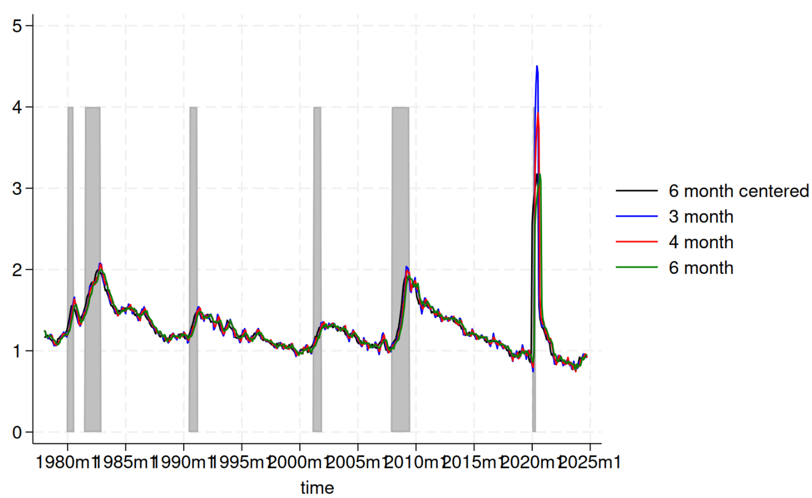

Figure 12: Comparison of moving-average filter for layoff series

A.2 DeNUNifying the Data

One common concern when linking individuals or household in the CPS data is that unemployment and non-participation are misclassified. In the following we will provide the main statistics for our data in the main text and deNUNified data. For the deNUNified data we remove all individuals which make one of the following labor market transitions: non-participation to unemployment to non-participation, or unemployment to non-participation to unemployment.

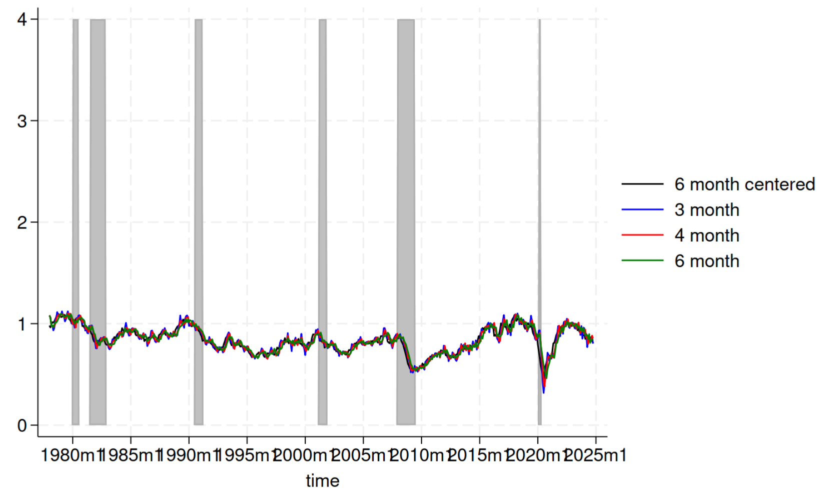

Figure 13: Comparison of moving-average filter for quit series

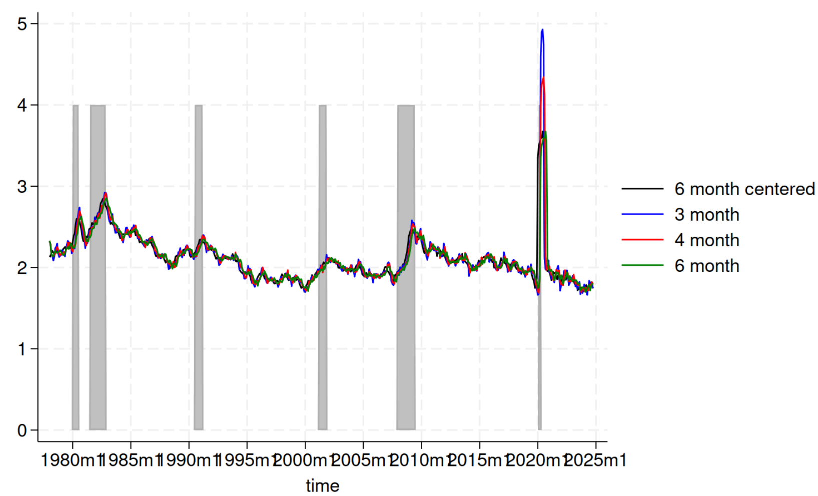

Figure 14: Comparison of moving-average filter for total separations

Table 16 shows that excluding these potentially misclassified transitions has no effect on the main statistics in this paper.

Permanent vs. Temporary Layoffs When we harmonize the data to compute the layoff rate, the possible answer choices of unemployed and non-participating individuals in the CPS can be grouped in two categories: layoffs from a temporary job or from a permanent job. The former category includes everyone who reports losing their job because a temporary, seasonal, or intermittent job ended. The latter includes all other job losers.

Figure 15 shows that the business cycle pattern of layoffs is driven by permanent layoffs and the majority of layoffs are from permanent jobs. Interestingly, layoffs from temporary are mildly

Statistic	Main Data	DeNUNified Data
	Averages	
Quits	0.84	0.84
Layoffs	1.27	1.27
Total Separations	2.11	2.11
Layoffs share N	0.40	0.40
Quit share N	0.85	0.85
EN	1.54	1.54
EU	1.08	1.08
	Correlation with Unemployment Rate	
EUQ	-0.1478	-0.1464
ENQ	-0.3897	-0.3904
EQ	-0.3929	-0.3917
EUL	0.6013	0.6011
ENL	0.4885	0.4872
EL	0.6117	0.6119
EN	0.0122	0.0119
EU	0.6210	0.6213
Layoff share N	-0.6222	-0.6239
Quit share N	-0.1253	-0.1302
Corr(EQ,EL)	-0.3019	-0.3007

Table 16: Comparison with deNUNified data

procyclical, i.e. decline during periods when unemployment is high. Because permanent layoffs are strongly countercyclical and temporary layoffs mildly procyclical, recessionary periods are characterized by an increase in the share of layoffs from a permanent job.

B Working-Age Population

This section provides the same figures and statistics as in the main text but for the working-age population, i.e. everyone in the United States who is 16 years or older and not currently institutionalized or an active member of the armed forces.

Statistic	Quits			Layoffs			Total sep.
	to U	to N	Total	to U	to N	Total	to (U+N)
Corr(x, y)	-0.0626	-0.2469	-0.1906	0.5083	0.3982	0.5775	0.4499
SD(x)/SD(y)	0.0300	0.1451	0.1648	0.2909	0.0894	0.3200	

Table 17: Business cycle correlations of each flow (x) with the unemployment rate (y)

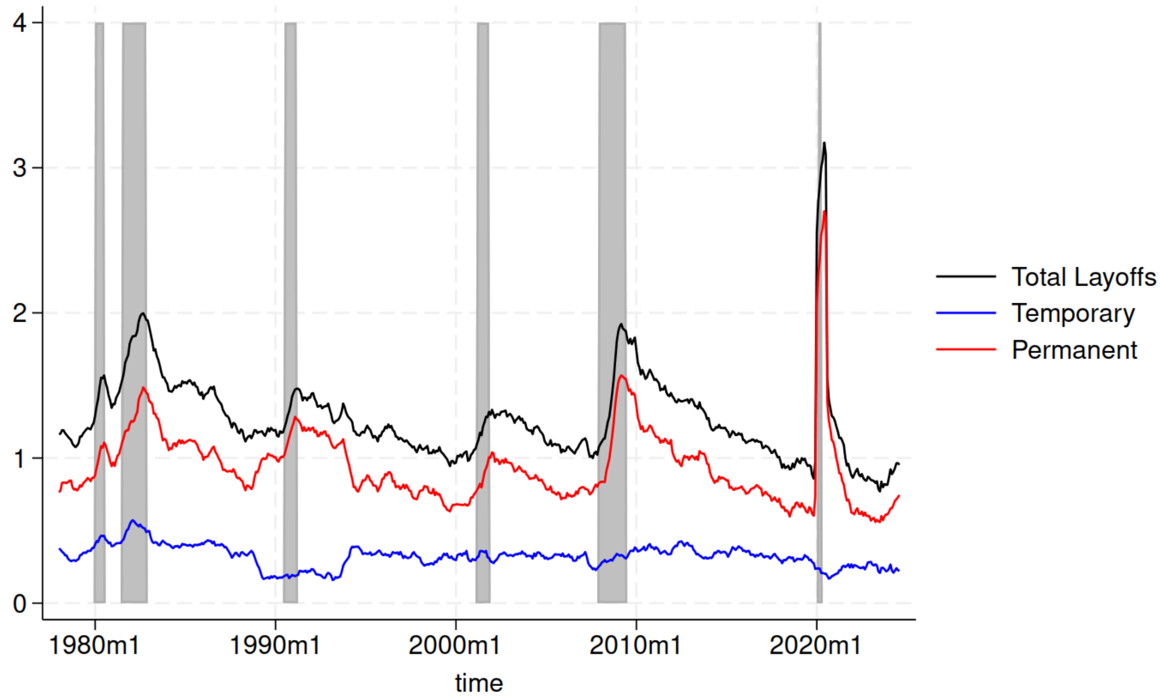

Figure 15: Layoffs from temporary vs permanent jobs

Statistic	
Corr(EQ, EL)	-0.0255
SD(EQ)/SD(EL)	0.5149

Table 18: Business cycle correlations of quits and layoffs

C Comparison to Jolts and Other Data

The Job Openings and Labor Turnover Survey (JOLTS) has been the primary source used to analyze quits and layoffs in the United States.[22] It is a monthly employer survey run by the Bureau of Labor Statistics (BLS). In this section, we compare our CPS quits and layoffs series with the corresponding JOLTS series.

JOLTS defines layoffs as "Involuntary separations initiated by the employer" and quits "Employees who left voluntarily. Exception: retirements or transfers to other locations are reported with Other Separations". Lastly, the JOLTS category "Other Separations" includes "retirements; transfers to other locations; deaths; or separations due to employee disability". Therefore, a quit in JOLTS is any voluntary separation with the exception of retirement, disability, death, or transfers

[22] Other complementary and timely data sources include the Survey of Consumer Finance (SCF) as for example in Koşar and Van der Klaauw (2023).

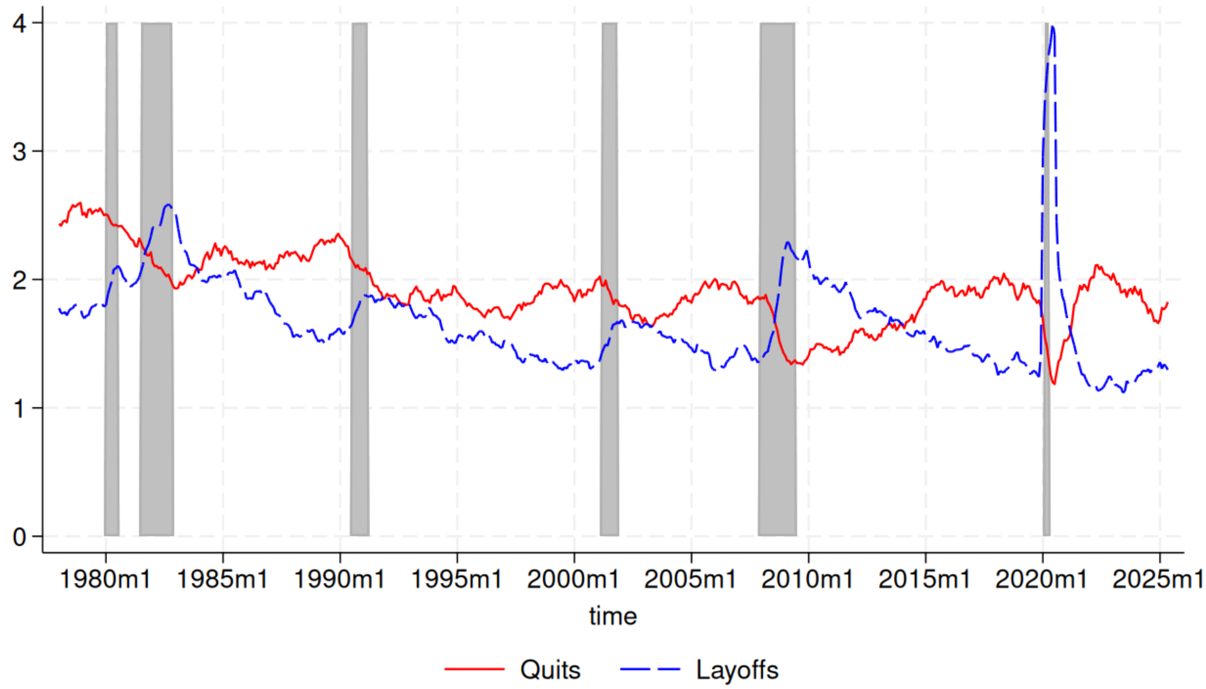

Figure 16: Quits and layoff

to other locations; and a layoff is any involuntary separation. It is important to note that JOLTS includes job-to-job quits and layoffs, whereas we can only observe the quits and layoff distinction for separations to non-employment[23] The JOLTS are also known to under count separations even when sampling weights are applied because they do not measure separations due to firm exit (Faberman (2005)). To remedy this, the disseminated JOLTS data are adjusted via a Monthly Alignment Method to produce stocks that are consistent with employment measured in the Current Employment Statistics (CES) (Cheng et al. (2009)).

In order to compare our data to JOLTS, we will restrict it accordingly. Layoffs are straightforward since we, similar to JOLTS, only consider individuals as laid off if they lost their job involuntarily. With regards to quits, we exclude all individuals who are retired[24] and disabled individuals are automatically excluded because they are not in the universe of individuals being asked the question of reason for non-participation. Death is also automatically excluded due to our linking strategy, because a dead person would not show up in the current month. Lastly, since we only consider separations into non-employment we do not have to worry about transfers to other locations. The earliest available from JOLTS is for January 2001, so restrict our series to start at

[23] Fujita et al. (Forthcoming) provide a series of employer to employer flows that does not distinguish quits and layoffs.

[24] By definition, they should not be asked the question in the CPS, but yet, there is a very small number in some months, which respond with retirement, and we exclude those

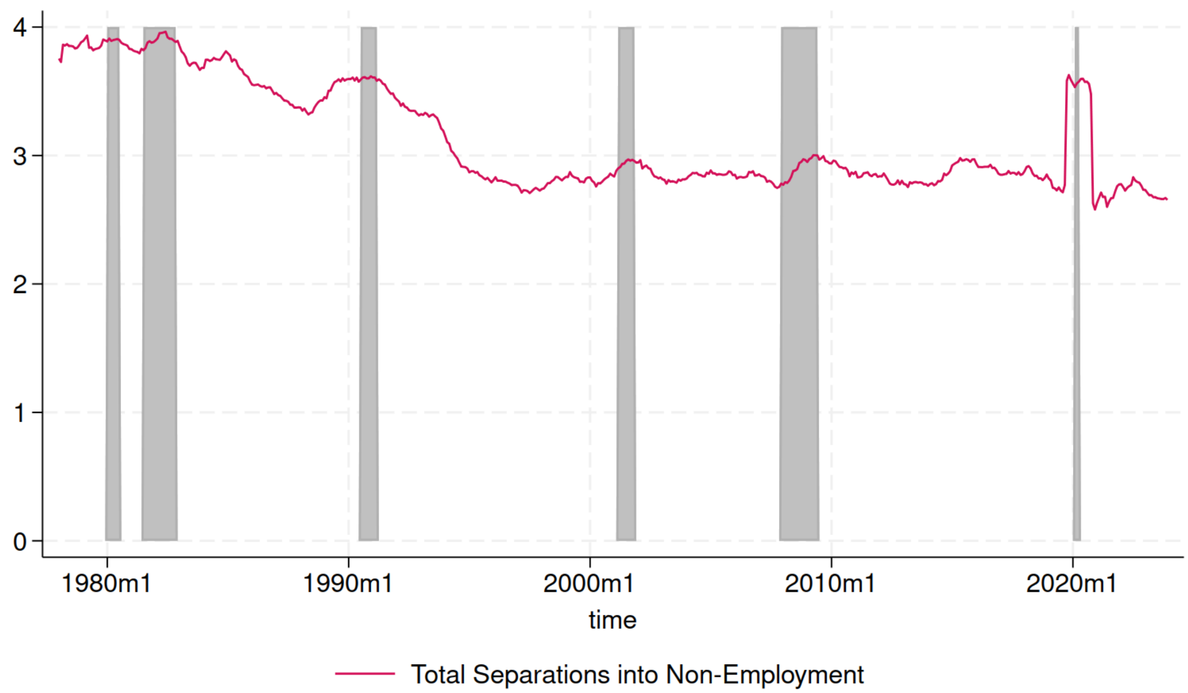

Figure 17: Total separations

the same date. Both series are seasonally-adjusted.

Figure 20 compares the JOLTS layoffs series with our layoffs series constructed using the CPS, including and excluding the pandemic recession. For every month in the sample, with the exception of the pandemic recession, the layoff rate computed using JOLTS data exceeds our layoff rate based on the CPS data. The correlation between the two series for the entire time period is 0.63. Notably, our layoffs series is significantly more responsive to fluctuations in the unemployment rate. The correlation of the CPS layoffs series with the unemployment rate is 0.50, whereas it is only 0.27 for the JOLTS layoffs series.

Comparison to Panel Study of Income Dynamics (PSID) The PSID is a long-running panel survey that has grown to over 9,000 families. While the smaller sample size limits the accuracy of business cycle analysis in the PSID relative to CPS or JOLTS, the comprehensiveness of the survey surpasses the other two sources. Using data from 2003-2019 we can study the reason each *job* an individual has ended (similar to JOLTS) and the labor market status of the individual after a job ends (as in CPS). The PSID is, for these reasons, an excellent check on the accuracy of our CPS classification and can reconcile some differences with JOLTS.

Table 19 shows that the split between quits and layoffs to non-employment that end up in non-participation is similar in the PSID as it is in our CPS sample. Over the same period, 12.9% of prime

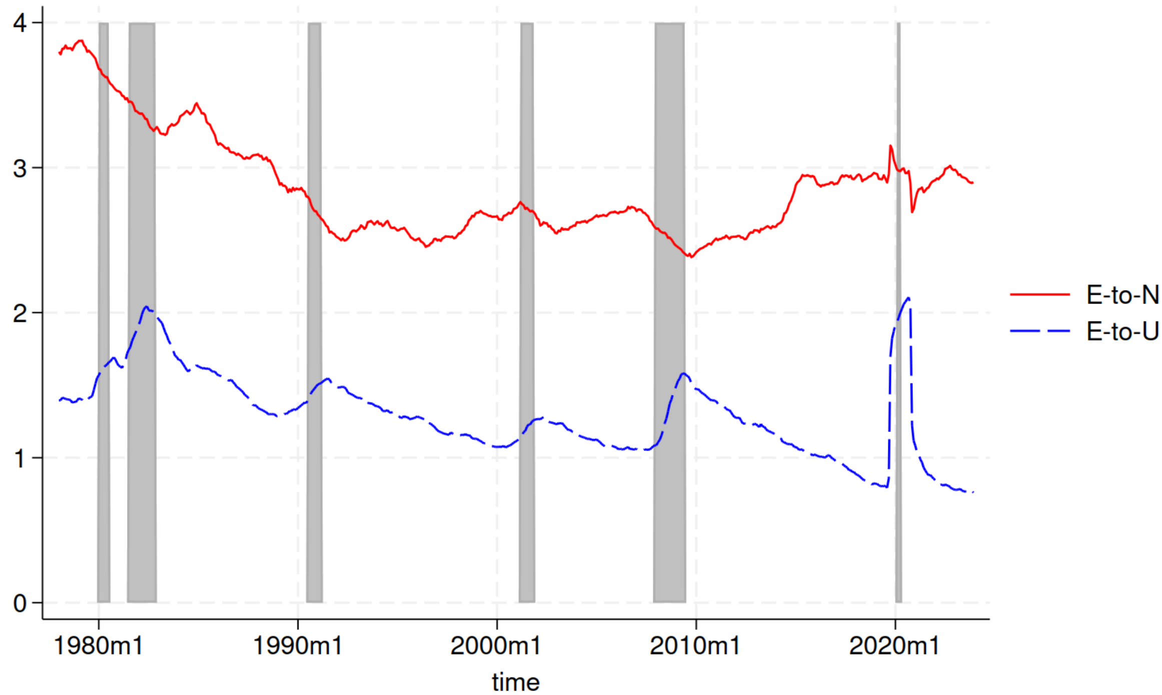

Figure 18: EN and EU flow rates

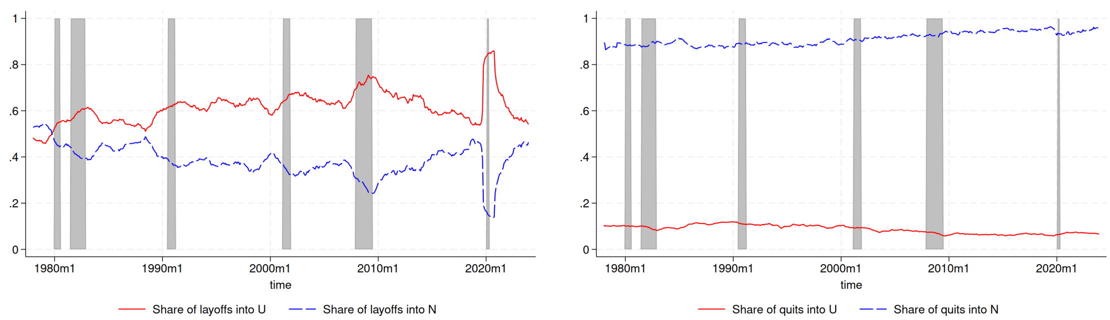

Figure 19: Share of quits and layoffs by destination

age quits to non-employment are classified as unemployment in the CPS compared to 11.8% in the PSID sample; and 64.8% of prime age layoffs to non-employment are to unemployment compared to 68.0% in the CPS sample.[25] This provides confidence that our classification of quits and layoffs is consistent with how workers describe the reason for a job ending in other popular surveys.

Table 19 also includes the separations we miss in the CPS. Separations directly to another employer or the termination of a single job held by a multiple job holder are included in quits and layoffs with a destination of "Employment". These types of separations make up 42.5% or 50.7%

[25] Quits to unemployment with an unemployment duration of over one year are dropped.

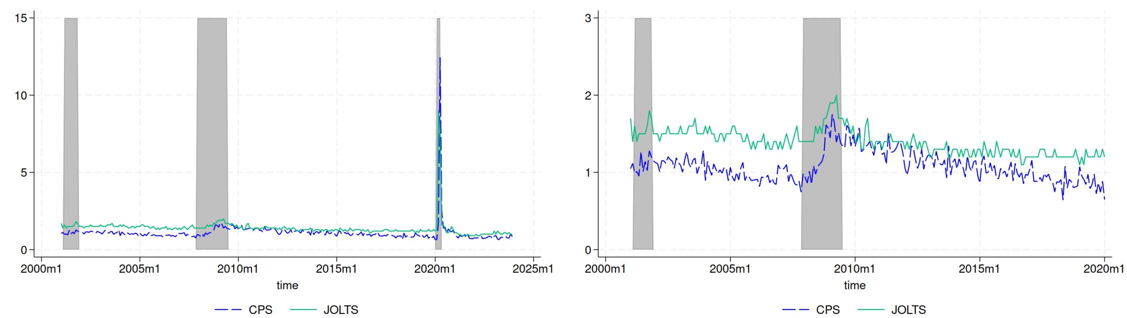

Figure 20: Layoffs, full series (left) and with the removal of 2020 + (right).

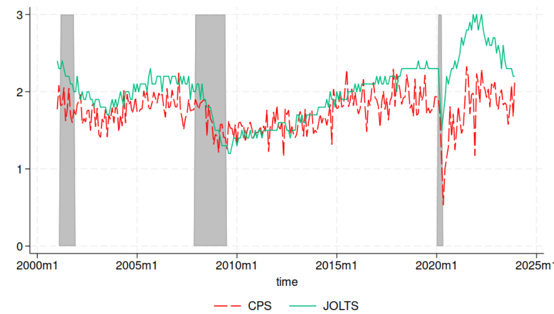

Figure 21: JOLTS total Quits and our adjusted CPS quit to nonemployment series

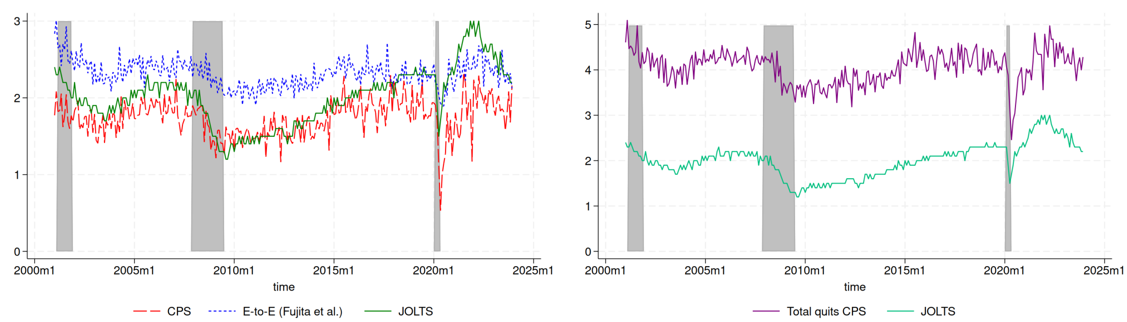

Figure 22: Quits: adjusted CPS quits to nonemployment, E-to-E flows Fujita et al. (Forthcoming), and JOLTS (Left); combined CPS quits plus E-to-E and JOLTS (right)

of all quits and 27.9% or 29.1% of all layoffs for all workers or prime age, respectively. During the recessionary years of 2008-10, the share of quits directly to a new employer falls to 30.1%; and the share of layoffs directly to a new employer falls to 21.6%. This backs our hypothesis that the CPS layoffs rise more during recession because layoffs to non-employment rise more than total layoffs, in part because the share of layoffs to non-employment increases. The analogous argument is supported for quits.

Table 20 broadens this comparison to a third survey, the Survey of Income and Program Partic-

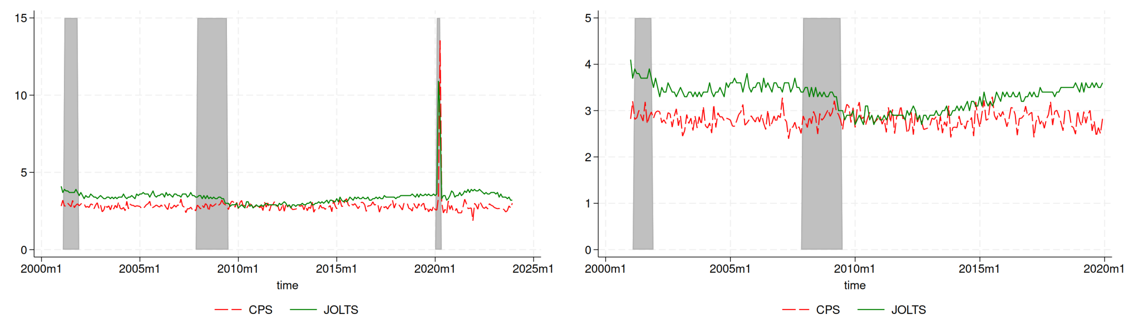

Figure 23: Total separations, full series (left) and with the removal of 2020 + (right).

Table 19: Panel Study of Income Dynamics 2003-2019

Share of Separations by Destination	All Workers		Prime Age	
Quits	All	to N or U	All	to N or U
Non Participation	0.540	0.912	0.452	0.882
Unemployment	0.052	0.088	0.060	0.118
Employment	0.425	n/a	0.507	n/a
Layoffs	All	to N or U	All	to N or U
Non Participation	0.273	0.349	0.248	0.320
Unemployment	0.509	0.651	0.527	0.680
Employment	0.279	n/a	0.291	n/a

ipation (SIPP, 2019–2023), and reports destination shares both for *total* flows (including job-to-job moves, observable in the panel surveys but not the CPS) and for flows to non-employment only. The qualitative picture is the same across the CPS, PSID, and SIPP: quits flow overwhelmingly to non-participation (the QN share of quits to non-employment is 0.87 in the CPS, 0.88 in the PSID, and 0.79 in the SIPP), while layoffs split much more toward unemployment. Including the job-to-job margin, about half of quits and a quarter of layoffs are direct employer-to-employer moves, and both of these direct shares fall in recessions (the recession columns), reinforcing the point made in the main text that the CPS layoff rate rises more in recessions because the share of layoffs flowing to non-employment, rather than to a new job, increases.

Finally, two distributional pieces of PSID evidence, referenced in the main text in Section 8, support reading the destination of a separation as a signal of a worker's distance to the participation margin. Figure 24 plots the distribution of residual earnings in the month before a separation, by transition type, where residuals are taken from a regression of earnings on education interacted with a cubic in age. Workers who quit to non-participation have the most leftward (lowest) distribution;

	Total Flows			to Non-Employment				
	PSID		SIPP	CPS		PSID		SIPP
	All	Recessions	All	All	Recessions	All	Recessions	All
QN	0.452	0.547	0.416	0.866	0.854	0.882	0.874	0.791
QU	0.060	0.079	0.110	0.133	0.146	0.118	0.126	0.209
QE	0.507	0.391	0.474					
Q	1	1	1	1	1	1	1	1
LN	0.248	0.199	0.241	0.359	0.310	0.320	0.242	0.389
LU	0.527	0.623	0.373	0.643	0.689	0.680	0.758	0.611
LE	0.225	0.179	0.387					
L	1	1	1	1	1	1	1	1

Table 20: Share flowing to destination: (N)on-participation, (U)nemployment, and directly to another (E)mployer, by type of flow: (Q)uit or (L)ayoff. PSID: 2002–2018; SIPP: 2019–2023.

those laid off into unemployment the most rightward (highest); and layoffs to non-participation lie in between. This is the surplus ordering the model assigns: selective quits are low-surplus workers leaving the labor force, while the more attached workers are those laid off who then search.

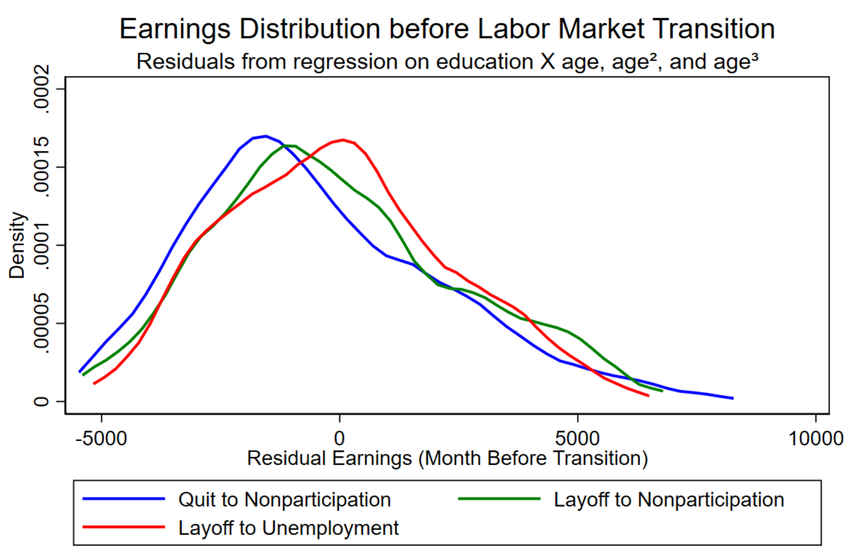

Figure 24: Distribution of residual earnings in the month before a labor market transition, by type of transition (PSID). Residuals are from a regression of earnings on education interacted with a cubic in age.

Figure 25 makes the same point with non-employment duration. Spells are shortest, and most concentrated at short durations, for layoffs to unemployment (LU); quits to non-participation (QN) and layoffs to non-participation (LN) carry the long right tails. Interpreting the persistence of non-employment as related to distance from participation, this confirms that destination-N separators are further from the margin and that our finer flow classification, which conditions on destination as well as on the reason for separation, adds information.

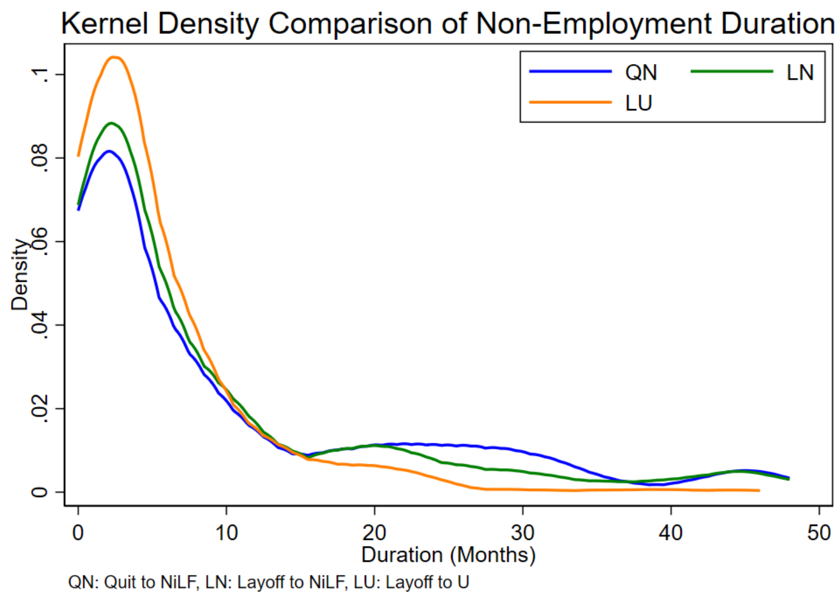

Figure 25: Kernel density of non-employment duration by transition type (PSID): quit to non-participation (QN), layoff to non-participation (LN), and layoff to unemployment (LU).

D Computational Appendix

This appendix details the algorithm summarized in the Computation and Estimation sections of the main text. The period is one month; all quantities are monthly.

D.1 State space and discretization

The individual state is (a, z, p, e): assets a, idiosyncratic productivity z, the participation-constraint indicator p, and (for the unemployed) UI eligibility e. Assets lie on a grid of $N_a = 48$ points spanning [0, 30] in units of the median monthly wage, with quadratic spacing ($a_i = 30\,(i/47)^2$) so points concentrate near the borrowing constraint where the consumption policy is most curved. Log productivity follows $\log z' = \rho_z \log z + \varepsilon$, $\varepsilon \sim N(0, \sigma_\varepsilon^2)$, discretized on $N_z = 15$ points by the method of Tauchen (1986): the grid spans ± 3 unconditional standard deviations and σ_ε^2 is set so the stationary variance of $\log z$ equals 0.20 at the estimated ρ_z. We carry five labor-market "statuses," namely employed (E), unemployed without and with UI eligibility (U_0, U_1), and free and participation-constrained non-participants (N_F, N_C), so that eligibility and the participation constraint are tracked as discrete states rather than approximated. The continuation-cost (selective-layoff) shock arrives i.i.d. with probability $\lambda_x(Z)$; the participation process moves workers between the free and constrained non-participation states with persistence (p_0, p_1); eligibility is lost at rate $\mu(Z)$.

D.2 Stationary equilibrium

For a candidate labor tax τ we solve the five value functions by value-function iteration to a sup-norm tolerance of 10^{-7}. Each backward step forms, for every status, the expected continuation value $\mathbb{E}_{z'|z}[\,C_s(a', z')\,]$ from next period's value functions (where C_s embeds the discrete separation, job-offer, eligibility, and participation lotteries described in Section 8) and then maximizes $u(c) - \theta\mathbf{1}\{\text{search}\} + \beta(1 - R_d)\,\mathbb{E}[C_s]$ over a' subject to the budget constraint with the consumption floor $\underline{c}$ imposed. The savings policy is monotone in a within each (s, z), which we exploit so the maximization is $O(N_a)$ rather than $O(N_a^2)$. The invariant distribution over (s, a, z) is computed by non-stochastic iteration of the implied transition operator to a tolerance of 10^{-11}; because the operator is applied exactly, the resulting moments carry *no* simulation noise, which is what lets the SMM objective be optimized by derivative-based and simplex methods without Monte-Carlo error. Alongside the distribution we iterate a tagged sub-distribution D^{sel} that records, in each non-employed cell, the mass whose current non-employment spell began with a *selective* separation (a selective quit or a selective layoff); the ratio D^{sel}/D delivers the stock-selectivity diagnostics in Table 9. Newborns (mass R_d) replace the deceased, enter non-employed holding average assets, draw z from its stationary distribution, and are counted as originating from non-participation. Finally, the labor tax is the fixed point of the period-by-period government budget, $\tau = (\text{UI outlays}+\text{floor outlays})/\text{wage bill}$, found by damped iteration. Prices come from the capital–effective-labor ratio $\kappa = K/L$: $r = r^\star A(Z)(\kappa/\kappa^\star)^{\alpha_y - 1}$ and $w = A(Z)(\kappa/\kappa^\star)^{\alpha_y}$, with K aggregate household savings and L aggregate efficiency units of employed labor, normalized so $w^\star = 1$ and $r^\star$ delivers a one-percent annual return.

D.3 Business cycles as a vector impulse

A recession is a single unanticipated MIT vector impulse, following Boppart et al. (2018). Rather than holding the cyclical parameters at their bad-state values for a fixed spell, we let the business cycle be one innovation that propagates with persistence: at the impact month every cyclical parameter $\Theta \in (A, \delta, \lambda_u, \lambda_n, \lambda_x, \mu)$ jumps to its bad-state value, and thereafter its deviation from the normal-times value decays geometrically at the common rate ρ_x,

$$\Theta_t = \Theta^G + (\Theta^B - \Theta^G)\,\rho_x^{t-1}, \qquad t = 1, 2, \ldots,$$

with $\Theta_0 = \Theta^G$. The economy is computed forward for a total of $T = 300$ months to approximate return to the steady state. We set $\rho_x = 0.95$ at the monthly frequency. The episodes of Section 11 use lower ρ_x values that match their own, faster persistence (0.92 for the 1981–82 recovery, 0.94 for the 2022–23 boom).[26] Given guesses for the paths of κ_t (hence prices) and the budget-clearing tax τ_t, we solve the household problem by backward induction from the steady-state value functions as the terminal condition, iterate the distribution forward from the steady-state distribution, and update (κ_t, τ_t) by damped fixed point (relaxation weight 0.5) until the path error falls below 2×10^{-4}; convergence is fast because prices and taxes move little over the cycle. A moment's recession "trough" is its largest absolute deviation from steady state along the impulse; under the geometric impulse the slow-moving stocks reach their troughs well after the one-month forcing, so we search a 180-month window. Recession averages are taken over the first 48 months. The volatility ratio $\text{std}(L)/\text{std}(Y)$ uses a 120-month window of log deviations, and the GDP decline is the peak-to-trough fall in $Y_t = A_t K_t^{\alpha_y} L_t^{1-\alpha_y}$ relative to its steady-state level.

D.4 How decomposition counterfactuals are constructed.

Three counterfactuals isolate the channels. (i) The *labor-supply* role freezes savings rules, discrete decision rules, and prices at their steady-state values along the recession path, while feeding in the *baseline* incidence of selective layoffs (the realized stay-vs-separate decisions), so the role captures only the quit, search, and participation responses and not changes in who is laid off. (ii) The *selection* role re-solves the recession's (δ^B, λ_x^B) so that the share of layoffs to non-participation stays at its steady-state value while the layoff-rate path matches the baseline (a two-equation root-find on window-average sN and window-average layoff rate), and compares the resulting transition to the baseline. (iii) The *one-shock-at-a-time* decomposition switches each cyclical parameter to its bad-state value alone, holding the others at normal, with agents understanding each path.

D.5 Estimation

We estimate by Simulated Method of Moments, minimizing a weighted sum of squared proportional deviations of the model moments from their targets, stacked over the normal and recession states,

[26] We experimented with a two-timescale impulse, a slowly-decaying aggregate (TFP) component together with fast-decaying separation/finding shocks, to recover the data's persistence without sacrificing the cross-sectional behavior. It does not help, for an instructive reason. Technology moves the unemployment rate almost not at all (Section 7), so a persistent TFP component does not lengthen the unemployment gap; the gap's persistence is governed by the persistence of the job-finding collapse. But a persistent finding collapse, while it preserves hoarding, makes non-participation a dead end and collapses the share of layoffs flowing there, whereas a fast-recovering finding rate restores that share but removes the reason to hoard. Unemployment persistence, hoarding, and the destination of layoffs are all governed by the same finding-rate clock and pull against one another; no decay structure delivers all three, and $\rho_x = 0.95$ is the scalar compromise.

plus the GDP-decline term. The weights are those given in the Estimation section: the employment-population ratio and unemployment rate carry the heaviest weights, the new separation series and attachment flows intermediate weights, and the basic gross flows the least. Because steady-state moments do not depend on the recession parameters, we estimate in two stages (stage A on the ten steady-state parameters and the normal-times moments, stage B on the five recession parameters and the recession deviations) and then polish jointly over all fifteen by multi-start Nelder–Mead, finishing by bisecting A^B so the trend-adjusted GDP decline is matched exactly with the other recession parameters re-optimized around it. The objective has the two basins discussed in Section 6.4; we report the attachment-consistent basin, in which recession UN falls and NU rises as in the data, and document all multi-start solutions in the replication materials.

D.6 Welfare

For each agent at the steady-state distribution we compute the permanent consumption-equivalent λ that equates the value of remaining at the steady state to the value of entering the recession at date zero, $\lambda = \exp\{(V^{ss} - V^{rec})\,(1 - \beta(1 - R_d))\} - 1$, where the annuitization factor $1 - \beta(1 - R_d)$ is the effective per-period discount including mortality. We report the mean and the breakdown by labor-market state at the recession's onset. The UI-extension experiment replaces $\mu^B = 0$ with the normal-times expiry rate μ^G and re-solves the transition; the selection-neutral welfare numbers use the counterfactual of the previous subsection.

E Identification, Standard Errors, and Sensitivity

This appendix formalizes the verbal argument of Section 6.4. Let $m(\theta)$ stack the normal-times and recession values of the eleven targeted moments and let θ collect the parameters governing labor supply (the search disutility θ and home production h) and the good- and bad-time frictions (the offer rates λ_u, λ_n, the random layoff rate δ, and the continuation-cost arrival λ_x). We compute the moment Jacobian $J = \partial m/\partial\theta$ by central finite differences at the estimated point, with steps scaled to each parameter.

Standard errors. Table 21 reports SMM standard errors from the sandwich formula $\text{Var}(\hat{\theta}) = (J'WJ)^{-1}J'W\,\Omega\,WJ\,(J'WJ)^{-1}$, where W is the estimation weighting matrix and Ω is a diagonal estimate of the targets' sampling variances, formed from the month-to-month variance of each CPS series over the steady-state and recession windows divided by the number of months. The parameters are precisely estimated. Most informative for this paper, the continuation-cost arrival λ_x and the participation parameters, the two ingredients that repair the model, carry large t-statistics, confirming that the separation-composition and attachment moments identify them rather than leaving them free.

Table 21: Identification: estimates, SMM standard errors, and t-statistics for the parameters governing labor supply (search disutility, home production) and the good- and bad-time frictions (offer rates, random layoff, continuation-cost shock). Standard errors use the sandwich formula with the estimation weighting matrix and diagonal target sampling variances from the monthly CPS series.

Parameter	Estimate	Std. error	t-stat
θ (disutility of search)	0.0953	0.0226	4.2
h (home production)	0.5119	0.0034	151.4
λ_u^G (offer rate U, good)	0.2733	0.0025	107.5
λ_n^G (offer rate N, good)	0.1119	0.0018	61.8
δ^G (random layoff, good)	0.0086	0.0002	34.7
λ_x^G (cont.-cost shock, good)	0.0857	0.0036	24.1
λ_u^B (offer rate U, bad)	0.2239	0.0073	30.5
λ_n^B (offer rate N, bad)	0.0610	0.0038	16.0
δ^B (random layoff, bad)	0.0152	0.0013	11.4
λ_x^B (cont.-cost shock, bad)	0.0921	0.0061	15.2

Sensitivity. Table 22 reports the local sensitivity matrix $S = (J'WJ)^{-1}J'W$ of Andrews et al. (2017), expressed as elasticities: entry (k, j) is the percent change in parameter k induced by a one-percent change in target j. The pattern confirms the verbal mapping of Section 6.4. The share of layoffs to non-participation loads most heavily on the continuation-cost arrival λ_x (and, through the estimated threshold, on the cost x): it is the moment that identifies selection. The quit rate and the non-participation re-entry rate NU load on the participation parameters (p_0, p_1), which split quits into their marginal and detached components. Home production h and the search disutility θ are identified jointly by the employment-population ratio, the quit level, and the destination of layoffs: the moments that locate the mass of workers near the participation margin.

Table 22: Local sensitivity of estimated parameters to the normal-times target moments: elasticities from the sensitivity matrix $S = (J'WJ)^{-1}J'W$. An entry is the percent change in the row parameter induced by a one-percent change in the column target.

Parameter	epop	urate	EU	EN	UE	NE	quit	layoff	share_layoff_N	UN	NU
θ (disutility of search)	13.24	-0.35	-0.00	-0.00	-0.38	-0.89	-0.00	-0.02	-0.73	-1.51	0.00
h (home production)	-1.24	-0.08	-0.00	0.00	0.09	-0.00	0.00	-0.00	-0.01	0.00	-0.00
λ_u^G (offer rate U, good)	1.27	0.07	0.00	0.00	0.76	0.07	0.00	0.00	0.07	-0.07	-0.00
λ_n^G (offer rate N, good)	0.14	-0.10	0.00	0.00	0.68	0.29	0.00	0.01	-0.14	0.26	-0.01
δ^G (random layoff, good)	-8.72	0.41	0.00	0.00	0.64	0.18	0.00	0.01	-0.76	0.02	-0.00
λ_x^G (cont.-cost shock, good)	-4.28	0.77	0.00	-0.00	0.03	-0.04	-0.00	0.00	1.49	-0.19	0.01
λ_u^B (offer rate U, bad)	-1.90	-0.17	-0.00	-0.00	0.14	-0.05	0.00	-0.00	-0.08	-0.03	0.00
λ_n^B (offer rate N, bad)	-5.96	-0.47	-0.00	0.00	-0.22	0.02	0.00	-0.00	-0.44	0.31	0.00
δ^B (random layoff, bad)	8.30	-0.23	-0.00	-0.00	-0.11	-0.18	-0.00	-0.01	-0.32	-0.19	0.00
λ_x^B (cont.-cost shock, bad)	20.48	1.14	0.00	-0.00	-0.87	0.03	-0.00	0.01	0.24	0.06	0.01

Recession accounting with standard deviations instead of steady-state to trough changes. Table 23 reports the sources of recessions decomposition in volatility terms, as the standard deviation each shock generates on its own.

Moment	All shocks	TFP	Job finding	Job loss	UI extension
Full model					
E/pop	1.03	0.00	0.45	0.54	0.09
U rate	14.29	0.05	5.30	8.10	2.11
UE rate	3.91	0.00	3.91	0.00	0.00
UN rate	7.08	0.35	7.56	2.73	8.03
NU rate	10.26	0.18	11.87	0.71	0.89
Quit rate	8.58	0.00	12.10	0.39	0.06
Layoff rate	8.41	0.00	1.01	9.77	0.57
Share of layoffs to N	8.04	0.00	1.47	4.84	1.17
Standard model					
E/pop	0.91	0.25	0.08	0.50	0.07
U rate	11.02	1.17	1.95	6.06	3.60
UE rate	0.60	0.00	0.60	0.00	0.00
UN rate	16.59	2.16	2.56	0.36	16.54
NU rate	8.61	6.75	15.03	3.85	15.76
Quit rate	1.80	2.48	4.25	0.73	0.54
Layoff rate	6.69	0.00	0.00	6.69	0.00
Share of layoffs to N	4.46	3.31	3.42	0.21	2.33

Table 23: What drives recessions, in volatility terms. Each entry is the standard deviation of the HP-filtered log of the moment (percent) generated along the impulse when only that shock is active.

Table 24 reports the labor-supply and selection roles in volatility terms.

	Full model			Standard model	
Moment	Baseline	Labor supply	Selection	Baseline	Labor supply
E/pop	1.03	-0.39	-0.31	1.16	+0.44
U rate	14.29	+1.86	+2.66	12.22	+0.57
UN rate	7.08	-1.62	+2.88	15.98	+0.78
NU rate	10.26	+9.18	+4.50	11.72	+6.51
Quit rate	8.58	+7.17	+3.15	7.98	+7.64
Layoff rate	8.41	+0.74	-0.24	8.80	+0.00
Share of layoffs to N	8.04	-0.44	+7.28	2.97	+2.75
Want-to-work share	4.96	-0.55	-0.58	4.34	+1.20

Table 24: **Roles of labor supply and selection, in volatility terms.** "Baseline" is the standard deviation of the HP-filtered log of the moment (percent) along the impulse; "Labor supply" and "Selection" are how much of that volatility each channel adds, measured as the baseline volatility minus the volatility with that channel switched off.