A Theory of How Workers Keep Up With Inflation*

Hassan Afrouzi
Columbia University

Andres Blanco
FRB Atlanta

Andres Drenik
UT Austin

Erik Hurst
Chicago Booth

December 2024

Abstract

In this paper, we develop a model that combines elements of modern macro labor theories with nominal wage rigidities to study the consequences of unexpected inflation on the labor market. The slow and costly adjustment of real wages within a match after a burst of inflation incentivizes workers to engage in job-to-job transitions. Such dynamics after a surge in inflation lead to a rise in aggregate vacancies relative to unemployment, associating a seemingly *tight* labor market with *lower* average real wages. Calibrating with pre-2020 data, we show the model can simultaneously match the trends in worker flows and wage changes during the 2021-2024 period. Using historical data, we further show that prior periods of high inflation were also associated with an increase in vacancies and an upward shift in the Beveridge curve. Finally, we show that other "hot labor market" theories that can cause an increase in the aggregate vacancy-to-unemployment rate have implications that are inconsistent with the worker flows and wage dynamics observed during the recent inflationary period. Collectively, our calibrated model implies that the recent inflation in the United States, all else equal, reduced the welfare of workers through real wage declines and other costly actions, providing a model-driven reason why workers report they dislike inflation.

JEL Codes: E24, E31, J31, J63
Key Words: Inflation, Vacancies, Job-to-Job Flows, Beveridge Curve, Wage Growth

*We thank Yash Singh for his outstanding research assistance and Guido Lorenzoni, Gabe Chodorow-Reich and Joe Hazell for providing us with helpful comments. We also thank seminar participants at Chicago, Columbia, Harvard, LSE, Michigan, UCLA Anderson, the Federal Reserve Board, the Federal Reserve Bank of Atlanta, the 2024 Society of Economic Dynamics conference, the 2024 NBER Macroeconomics over Time and Space conference, the 2024 NBER Monetary Economics Fall meeting, and the 2024 St. Louis Fed-Washington University Macro Labor conference for helpful comments. The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve Board. Email: hassan.afrouzi@columbia.edu; julio.blanco@atl.frb.org; andres.drenik@austin.utexas.edu; and erik.hurst@chicagobooth.edu.

1 Introduction

Decades of low and stable inflation in the U.S. ended with the inflation spike of 2021. Whereas inflation had hovered annually at around 2.2 percent between 2000 and 2019, prices rose by over 14 percent cumulatively between April 2021 and May 2023. Yet, unemployment remained low and relatively stable during this period. Instead, vacancy postings shot up, and labor market tightness, measured by the aggregate vacancy-to-unemployment rate (V/U), reached historically high levels, as shown in Panel A of Figure 1. High inflation, low unemployment, and a high vacancy-to-unemployment rate all pointed towards an economy that was "running hot" with too many firms chasing after too few workers, a narrative that was echoed by both policymakers and academics. In his post-FOMC press conference on November 2, 2022, Chairman Powell declared that "the broader picture is of an overheated labor market where demand substantially exceeds supply." Soon after, benchmark New Keynesian Phillips curves—which had not needed a role for vacancies since their inception—were adapted to explain how a tight labor market can cause a surge in inflation through higher vacancies, even with modest changes in unemployment.[1]

But was the labor market "hot" during this period? Not according to real wages, which fell sharply with the rise in inflation and continue to be below their pre-trends. As seen in Panel B of Figure 1, real wages for the median worker in June 2024, as measured by Atlanta Fed's Wage Tracker Index, were still about 4.4% below where they were predicted to be based on pre-2020 trends; during the 2016-2019 period, real wages for the median worker increased at a rate of 1.4% per year. Consistent with the declining real wages, survey evidence has documented that workers unambiguously perceive their welfare to have declined during the recent inflation period.[2] The juxtaposition of the seemingly "hot labor market" implied by the rising vacancy-to-unemployment rate with the decline in real wages questions the role of the tight labor market in raising prices during the recent period.

In this paper, we argue that the causation can go in the opposite direction, where unexpectedly high inflation drives a rise in the vacancy-to-unemployment rate, giving the appearance of a tight labor market while simultaneously generating a period of declining real wages. To formalize this idea, we combine modern models of aggregate labor market

[1] For a recent discussion of how an increase in the vacancy-to-unemployment ratio in particular, and labor market tightness in general, can cause higher inflation see, for example, Ball, Leigh, and Mishra (2022), Bernanke and Blanchard (2024), Benigno and Eggertsson (2023), and Lorenzoni and Werning (2023b).

[2] See Stantcheva (2024) and Afrouzi, Dietrich, Myrseth, Priftis, and Schoenle (2024).

Figure 1: Vacancy-to-Unemployment Rate, CPI and Real Wages Over Time

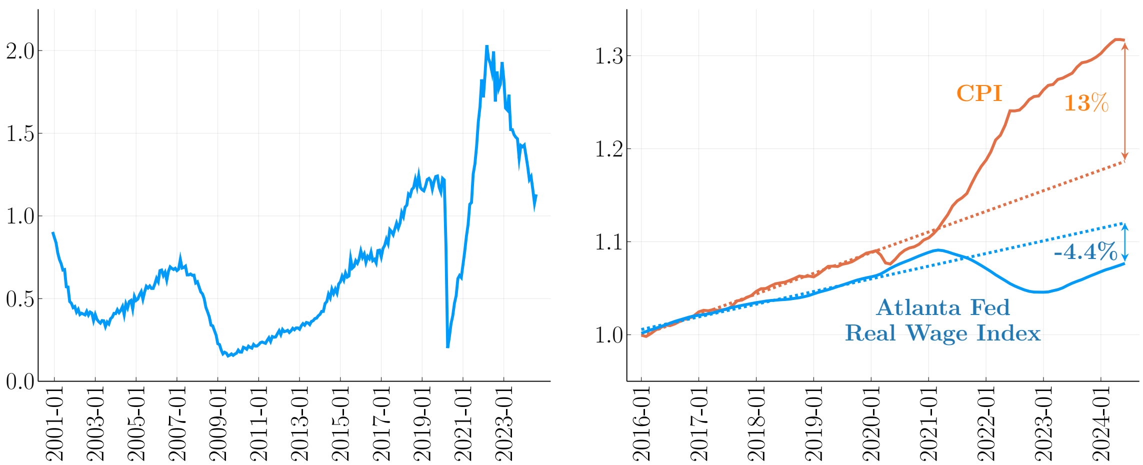

Panel A: Vacancy-to-Unemployment Rate: 2000-2024

Panel B: CPI and Real Wages: 2016-2024

Notes: Panel A shows the vacancy-to-unemployment rate from 2001M1 through 2024M8, where vacancies come from the JOLTS survey. Panel B shows the evolution of the CPI (red line) and the Atlanta Fed's Nominal Wage Index deflated by the CPI (blue line); we refer to the latter as the Atlanta Fed Real Wage Index. The dashed lines in the figure project the growth rate in each series from January 2016 and December 2019 over the entire sample period. See Section 2 for additional details on the series construction.

flows with nominal wage rigidities to explore the aggregate and distribution consequences of "inflation shocks", both theoretically and quantitatively, on labor market outcomes and worker well-being.[3] In particular, we incorporate nominal wage rigidities into a model with heterogeneous workers and frictional labor markets with many types of endogenous worker flows (quits, layoffs, and on-the-job search) to explore the effects of inflation on worker wages and welfare. In this environment, a burst of inflation, all else equal, reduces real wages on impact. In order to have their real wages keep up with inflation, workers can take steps to renegotiate their wage with their existing firm at a cost, engage in costly search for another job where they could contract over an updated real wage, or simply quit to unemployment if they find their eroded wages to be too low. The culmination of all these endogenous

[3]In the model, the "inflation shock" will be an exogenous increase in the price level. The goal of the paper is not to explain the causes of the recent inflation but, instead, to assess how inflation affects labor market dynamics, all else equal. However, in the last section of the paper, we discuss how prominent stories for the current inflation episode would effect labor market dynamics.

forces increases quits, job-to-job flows, and vacancies with modest effects on unemployment, resulting in a higher vacancy-to-unemployment ratio *caused by* unexpectedly lower real wages. We then validate the predictions of our model using historical U.S. data between 1950 and 2019 to show that prior inflationary periods also resulted in a high vacancy-to-unemployment rate and an upward shift in the Beveridge Curve, all else equal. Finally, we conclude by showing that other types of hot labor market shocks that can generate a large increase in the vacancy-to-unemployment rate yield implications that are inconsistent with observed wage dynamics and other labor market flows during the recent U.S. inflationary period.

We begin the paper by documenting a series of facts about the U.S. labor market during the recent inflation period. We use these facts to both motivate the elements we include in the model and to evaluate the success of our model by its ability to match these facts. Using data from the *Job Openings and Labor Turnover Survey* (JOLTS), we show that, starting in April 2021, the layoff rate plummeted by roughly 20%, and the quit and vacancy rates jumped by roughly 20% and 50%, respectively, relative to the 2016-2019 period. We show that the monthly quit rate and the monthly vacancy rate were very closely correlated with the monthly inflation rate during this period. The quit and vacancy rates were the highest when the inflation rate was the highest. Importantly, we also show that there was no change in the job-finding rate of the unemployed during this period. Using data from the *Current Population Survey*, we document that E-E flows jumped for individuals throughout the income distribution during the inflation period with the increase being largest for lower-wage workers. Likewise, we show that real wages fell for workers throughout the income distribution with the declines being largest for higher-wage workers. Finally, we show that compared to the pre-period, nominal wage growth grew significantly more for job-changers than job-stayers.

Motivated by these patterns, we develop a modern macro-labor search model with sticky wages, where wage renegotiation incurs costs and renders wage changes infrequent, consistent with the observed data. Along with frictional adjustment of wages, we assume firms and workers lack commitment, so that at any given time, based on their current states, workers decide whether to renegotiate their wage, quit into unemployment, or search for a new job, while firms determine whether to lay workers off or retain them.[4] Our model postulates two

[4]Recently, Blanco, Drenik, Moser, and Zaratiegui (2024) illustrates how adding in nominal wage rigidities and two-sided lack of commitment into otherwise standard modern models of frictional labor markets can generate inefficient job separations. Our framework shares many of the insights about inefficient separations when nominal wages are sticky. Such endogenous quits, layoffs, and wage renegotiations with fixed costs introduce the key ideas of models of inaction in the pricing literature (e.g., Barro, 1972, Sheshinski and Weiss,

main channels for employed workers to overcome the stickiness of their nominal wages. First, for workers who remain with their employer, their wages can adjust to a level that is not greater than the long-run target inflation rate with no additional utility cost; we model this adjustment process with an exogenous Poisson arrival rate. Additionally, workers can also pay a randomly drawn fixed cost, that is finite with some Poisson arrival rate, which allows them to renegotiate their real wages. Second, we assume that the wages of new hires are fully flexible, meaning that workers can also adjust their nominal wages by searching on the job (also at a cost) and potentially moving to a match with a new employer. Job search is frictional and directed on the part of both workers and an infinite mass of homogeneous firms. Finally, we assume that the flow benefits to the non-employed are set in real terms.

To examine how different workers are affected by an unexpected temporary burst of inflation, the model includes heterogeneous worker types who differ in their latent productivity. By incorporating worker heterogeneity, we can assess how well our novel framework can explain both the time series and the cross-sectional response of worker flows and wages to aggregate changes in the inflation rate. In addition to ex-ante heterogeneity, the productivity of the employed (unemployed) evolves over time according to i.i.d. Brownian motions with positive (negative) drift. We also allow the worker's flow benefits of non-employment to flexibly scale with worker productivity. For example, we allow for the possibility that the wages of low-productivity workers are, on average, closer to or farther away from the outside option of non-employment; we later let the data speak to either of these possibilities. In a similar vein, we also allow vacancy posting costs to flexibly scale with worker productivity. This allows for the possibility that it is either more or less expensive to hire a high-productivity worker relative to a low-productivity worker. As discussed below, we discipline these scaling factors with micro data on differences in job-finding rates and job-to-job flows across the wage distribution during the pre-inflation period. We estimate that the value of non-employment is relatively higher and the cost of posting a vacancy is relatively lower for low-productivity workers implying they are more responsive to the inflationary shocks.

On the methodological front, the model requires solving endogenous decisions between matched workers and firms. Due to renegotiation costs and lack of commitment, we cannot rely on the usual equivalence to a planner's problem that maximizes the surplus of the match on behalf of firms and workers. Instead, we use a Markov Perfect Equilibrium concept

1977) into a modern macro labor model with search.

in continuous time to characterize both firms' and workers' decisions. Workers' strategies consist of which market to enter while unemployed, and once within a match, when to negotiate, when to quit, or when to search for a new job. Firms' strategies are when to layoff their employed workers. A technical contribution of our paper is to recast the strategic interaction between heterogeneous firms and employed workers as a stochastic non-zero sum game—since the surplus of a match is non-zero—with stopping times in continuous time. This approach characterizes the equilibrium conditions as two Hamilton-Jacobi-Bellman Variational Inequalities (HJBVIs) describing optimal policies and value functions, and allows us to use efficient numerical methods to solve for the equilibrium.

We use a variety of microdata sources in the years prior to 2020 to calibrate the key labor market parameters of our model. In particular, we use administrative payroll data on the frequency of wage changes and the distribution of the size of wage changes to calibrate the parameters governing nominal wage rigidities. The calibrated model matches additional non-targeted moments. In particular, our calibrated model implies that the average wage markdown is larger for higher-productivity workers relative to lower-productivity workers. This is because low-wage workers are more elastic to labor market shocks, giving firms less market power over these workers. The model prediction that wage markdowns are larger for high-wage workers is consistent with the empirical estimates using Danish microdata in Chan, Mattana, Salgado, and Xu (2023) and Norwegian microdata in Volpe (2024).

Using the calibrated model, we find that an unexpected temporary inflation shock of the size comparable to the inflation experienced in the U.S. during the 2021-2023 period matches key patterns observed in the labor market at that time. First, our model generates a rise in job search on the part of workers after a burst of inflation consistent with a variety of survey evidence collected during the recent inflationary period.[5] Second, this additional on the job search on the part of employed workers and leads to a rise in vacancies targeted towards employed workers, while vacancies targeted towards unemployed workers remain relatively stable. As a result, our model generates an *increase* in the aggregate vacancy-to-unemployment rate along with a *decline* in real wages, similar to the patterns observed in Figure 1. Third, and relatedly, our model also matches the upward shift in the Beveridge Curve observed in the U.S. economy during the last few years. In particular, our model

[5] Pilossoph and Ryngaert (2023) collected data showing the link between inflation expectations and labor search behavior during the recent inflation period. Likewise, Stantcheva (2024) documents that about 10% of workers reported actively searching for a new job as a result of the recent inflation.

predicts a large jump in vacancies due to increased E-E churn caused by rising inflation with little change in the aggregate unemployment rate.[6]

Moreover, we find that inflation reduced the average welfare of workers in all deciles of the income distribution. The losses were greatest, however, for higher wage workers who are relatively more inelastic; workers in the bottom decile, the median decile, and the top decile of the wage distribution experienced welfare losses from the current inflation of roughly 75%, 85%, and 110% of monthly real income, respectively. Additionally, even though we calibrate the model to data prior to 2020, we show that only feeding an appropriately sized inflation shock into our model generates time series patterns of firm layoffs, worker flows, vacancies, and wages that match closely the actual data. The model also shows that firms are better off from the burst of inflation because their market power increase; this finding is consistent with the historically high corporate profits to GDP ratio experienced by US firms during the 2021-2023 period. Lastly, we show that the majority of the welfare losses to workers come from the real wage declines. However, we also show that the increased search and renegotiation costs incurred by workers to have their real wages keep up with inflation further reduce worker welfare beyond the initial real wage declines. We also show that a reduced layoff margin creates welfare gains for workers that approximately offset the increased search and renegotiation costs.

After presenting the results of our model, we use historical data to show that high inflation rates systematically increase the vacancy-to-unemployment ratio and result in upward shifts in the Beveridge Curve. Using Barnichon (2010)'s unified vacancy series which combines data from the Conference Board's Help Wanted Index and JOLTS, we identify nine periods where the vacancy-to-unemployment rate substantially exceeded its long-run average. Four of those periods were associated with very high inflation and the unemployment rate was either at a relatively high level or, at least, non-decreasing: those periods were in the early 1950s, the mid-1970s, the late 1970s, and the current post-COVID period. All of these periods were marked by large negative aggregate supply shocks that contributed to the high inflation. The other periods of high vacancy-to-unemployment rates had relatively low inflation and a sharply declining unemployment rate consistent with traditional "hot labor-market" stories. We show that the vacancy rate and the vacancy-to-unemployment rate both systematically

[6]To that end, our paper provides additional supporting evidence for the mechanism highlighted in Cheremukhin and Restrepo-Echavarria (2023) which argues that the shape of the Beveridge Curve depends on the extent to which outstanding vacancies are filled with E-E transitions as opposed to U-E transitions.

increased when inflation was high during the 1950-2019 period conditional on the aggregate unemployment rate. We also discuss how similar increases and vacancies occurred during the high inflation in Argentina during the early 2000s. Collectively, these results provide additional empirical support for our theory.

Finally, we end the paper by discussing how the predictions of our model are distinct from the predictions of alternate "hot labor markets" theories that can generate an increase in the vacancy-to-unemployment ratio. We use our model to explore how (i) an increase in aggregate productivity, (ii) a decline in the household discount rate, (iii) a decline in the value of non-employment, and (iv) a decline in the vacancy posting cost can all generate an increase in the vacancy-to-unemployment rate; we calibrate the size of the shocks so that they replicate the increase in the vacancy-to-unemployment predicted by our baseline model. We then show that these other shocks fail to replicate other labor market dynamics that were observed during the recent inflation period. In particular, these theories struggle to match the decline in real wages, the higher real wage growth of job-changers relative to job-stayers, the rise in job-to-job flows, the fall in the layoff rate, and/or the relatively constant job-finding rate of the unemployed.

These findings suggest that academics and policymakers should be cautious about viewing the rise in the V/U rate as a sign of a hot labor market during inflationary periods without holistically looking at other labor market indicators. The causation from inflation to labor market dynamics implied in our model can make it appear that the labor market was hot (by increasing V/U rate) without underlying forces that would increase workers' real wages as in the hot labor market stories. We conclude the paper with a discussion of how our model is consistent with the fact that a combination of recent supply chain disruptions, energy price increases, and increased aggregate demand from Pandemic-related policies put upward pressure on prices but had relatively offsetting effects on labor demand.[7]

As discussed above, a key implication of our model is that accounting for the role of vacancies targeted towards employed vs. unemployed workers is key for understanding the recent rise in the aggregate vacancy-to-unemploymend ratio and the shift in the Beveride curve. In that sense, our paper is related to Moscarini and Postel-Vinay (2023), which introduces

[7] Lorenzoni and Werning (2023a,b) and Bernanke and Blanchard (2024) highlight the importance of commodity price increases, supply disruptions, and sectoral reallocation in explaining the recent rise in the U.S. price level which should reduce labor demand and put downward pressure on the vacancy-to-unemployment rate.

on-the-job search into a monetary DSGE New-Keynesian model and shows that the ratio of employer-to-employer transitions to unemployment-to-employment transitions (EE/UE) serves as a key predictor of inflationary pressures. Complementary to their mechanism, our paper demonstrates that inflation itself can alter the pattern of job-to-job transitions and vacancy creation, leading to shifts in the Beveridge Curve. Together, these mechanisms highlight the importance of distinguishing between job-to-job transitions and transitions from unemployment when analyzing the causes and effects of inflationary episodes.

Our work is also related to a set of recent papers showing how worker well-being is affected by recent inflation. First, our model provides a theoretical rationale for the survey findings in Stantcheva (2024) regarding declining worker welfare, while at the same time providing an explanation for how increasing job-to-job churn can cause a rise in the vacancy-to-unemployment ratio making it appear that the labor market was running hot. Second, Hajdini, Knotek, Leer, Pedemonte, Rich, and Schoenle (2022), Pilossoph and Ryngaert (2023), and Pilossoph, Ryngaert, and Wedewer (2024) all highlight how increased inflation can result in workers searching more for another job. Both Hajdini, Knotek, Leer, Pedemonte, Rich, and Schoenle (2022) and Pilossoph and Ryngaert (2023) use survey data to show that workers with higher inflation expectations increase job search effort. Separately, Guerreiro, Hazell, Lian, and Patterson (2024) fielded a novel survey in early 2024 asking respondents about whether they took costly actions—asking their boss for a raise, partaking in union activity, or soliciting external job offers—in response to the recent inflation activity. They find that about one-fifth of all workers engaged in costly actions to raise their wages during the recent inflation period. Our paper shows how jointly including these costly actions in a macroeconomic model of the labor market is necessary to replicate the labor market flows and wage dynamics observed during the recent inflation period.

2 Wages & Labor Market Flows in the Recent U.S. Inflation Period

We refer to the recent "*inflation period*" in the United States as beginning in April 2021 and extending through May 2023; for each month during this period, the year-over-year CPI inflation rate exceeded 4%. The cumulative price level increase exceeded 14% during this 26-month period. As a way of comparison, the inflation rate in the United States averaged about 2% per year during the 2000-2019 period and averaged 3.3% during the "*post-inflation*" period of May 2023 through June 2024.

In this section, we document a set of facts about how labor market flows and wages evolved during the recent inflation period within the United States both in the aggregate and across different income groups. Collectively, these patterns motivate the setup of our model described in the next section. In later sections, we evaluate the success of our model by its ability to match the broad time series patterns documented below.

2.1. Quits, Layoffs, and Vacancies During the Inflation Period

Figure 2 shows the trends in the monthly layoff rate, quit rate, and vacancy rate for the United States between 2016 and 2024 using data from the BLS's *Job Openings and Labor Turnover Survey* (JOLTS). The JOLTS dataset provides a snapshot of worker hiring and separation flows for a nationally representative sample of non-farm business and government employers during a given month.[8]

Layoff Rate. Panel A of Figure 2 shows the time series trend in the layoff rate prior to, during, and after the inflation period.[9] Between January 2016 and December 2019 (which we refer to throughout this section as the "*pre-period*"), the average layoff rate was fairly constant at about 1.22% per month. However, throughout the inflation period, the monthly layoff rate fell sharply to about 0.98% per month; during this period, the layoff rate was at its lowest level since the JOLTS data started in 2000. Relative to the pre-period, firms terminated workers at a much lower rate during the inflation period.

Quit Rate. Panel B of Figure 2 shows the time series trend in the quit rate during the 2016-2024 period. From 2016 through 2019, the quit rate averaged about 2.2% per month. During the inflation period, the quit rate jumped to an average of about 2.7% per month. The time series path of the quit rate followed closely the time series path of inflation; for example, both the inflation rate and the quit rate peaked in the second quarter of 2022. By early 2024, both the quit rate and the inflation rate had almost returned to their 2016-2019 levels.

[8] A detailed discussion of all data used in this section can be found in the Online Appendix. For this figure, we downloaded the data directly from the United State's *Bureau of Labor Statistics* (BLS) (`https://www.bls.gov/jlt/data.htm`). A discussion of how JOLTS measures the various labor market flows is also provided in the Online Appendix.

[9] To make the graph easier to read, we excluded the historic spike in the layoff rate during the beginning of the COVID recession from the figure. In March and April of 2020, the layoff rate jumped to 9.0% and 7.0%, respectively.

Figure 2: Layoff Rate, Quit Rate and Job Opening Rate 2016-2024, JOLTS Data

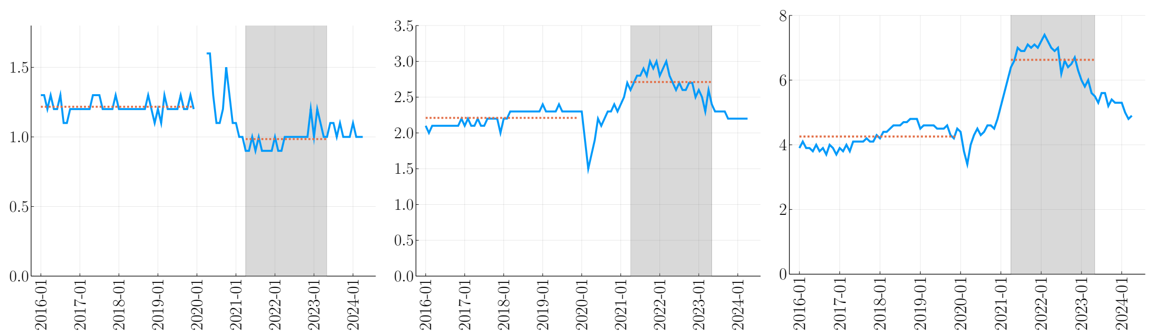

Panel A: Layoff Rate	Panel B: Quit Rate	Panel C: Vacancy Rate

Notes: Figure shows the layoff rate, the quit rate, and the vacancy rate for the U.S. economy from January 2016 through May 2024 using the BLS's JOLTS data. The shaded years highlight the inflation period. The dashed red lines show the average of the series during the 2016-2019 pre-period and then separately during the inflation period.

Vacancy Rate. The *vacancy rate* (or job-opening rate) is constructed as the number of open positions on the last business day of the month divided by the sum of employment and vacancies on the last day of the month. Panel C of Figure 2 shows the time series patterns of the vacancy rate, which closely follows the time series patterns of the quit rate; firms often post a vacancy to replace workers who quit. The average monthly vacancy rate jumped from 4.25% per month during the 2016-2019 period to 6.65% per month during the inflation period. Again, the time path of the vacancy rate tracked closely the time path of inflation during the 2021 to 2024 period.[10]

2.2. Worker Flows During the Inflation Period

The quit rate from the JOLTS data shown above captures workers who left the firm by either (i) flowing into unemployment before starting to look for another job (a voluntary "E-U" flow), (ii) directly transitioning to another firm (an "E-E" flow), or (iii) leaving the labor force (an "E-N" flow). Ellieroth and Michaud (2024) document that quits to non-employment did not increase during the 2021-2023 relative to the 2016-2019 pre-period; Appendix Figure B.4 also shows there was no increase in the flow of employed workers into unemployment during the 2021-2024 period. This suggests that the increasing quit rate from the JOLTS

[10] Online Appendix Figure A.1 shows that there is a very tight relationship between the monthly inflation rate and both the monthly quit rate (Panel A) and the monthly vacancy rate (Panel B) during the 2016-2024 period.

data was driven solely by job-to-job transitions. In this sub-section, we use data from the *Current Population Survey* (CPS) to further highlight that the increase in quits from the JOLTS data was driven by an increase in job-to-job flows and not driven by an increase of workers into non-employment.

Table 1 shows the average employment rate for different demographic groups in both the January 2016 to December 2019 pre-period and then again for the April 2021 to May 2023 inflation period. As seen from the table, the employment rate was essentially unchanged between the periods for the four demographic groups.[11] Likewise, the measured unemployment rate in 2018-2019 was essentially the same as the unemployment rate during the mid-2021 to mid-2023 period at roughly 3.7%. These results show that the recent inflation period was not associated with either substantively increasing employment rates nor with substantive changes in unemployment rates relative to the pre-pandemic period.

Table 1: Employment to Population Ratio Over Time, 25-55 Year Olds

Education	2016M1-2019M12	2021M4-2023M5
Men: Less than Bachelors	0.820	0.810
Men: Bachelors or More	0.920	0.918
Women: Less than Bachelors	0.664	0.660
Women: Bachelors or More	0.814	0.826
All	0.788	0.792

Notes: The first four rows of the table show the average employment rate for men and women with less than a Bachelor's degree and men and women with a Bachelor's degree or more in different time periods. The last row shows the average employment rate pooling men and women of all education levels. Column 1 shows the average employment rate during the January 2016 to December 2019 pre-period while Column 2 shows the average employment rate during the April 2021 to May 2023 inflation period. The sample focuses on those aged 25-55 from the monthly CPS files.

Panel A of Figure 3 shows the time series of the monthly E-E rate during the 2016-2024 period averaged at the quarterly frequency.[12] In the 2016-2019 period, the average E-E rate was about 2.30% per month. During the 26-month inflation period, the E-E rate jumped to

[11] Women with a bachelor's degree experienced a slight increase in their employment rate over this time period but this is a continuation of a trend that pre-dated 2020.

[12] For this analysis, we follow the procedure of Fujita, Moscarini, and Postel-Vinay (2024) to measure E-E flows in the CPS. See the Online Appendix for additional details.

an average of about 2.42% per month (p-value of difference < 0.01). In mid-2022, the E-E rate peaked at about 2.55% month. The CPS data complement the JOLTS data by showing that the increasing quit rate is accompanied by an increase in employer-to-employer transitions. Moreover, we show in Appendix Table B.1 that the E-E rate increased substantively for both lower- and higher-educated workers during the recent inflation period; the increase was larger for lower-educated workers.

Figure 3: E-E and U-E Flows 2016-2024, CPS Data

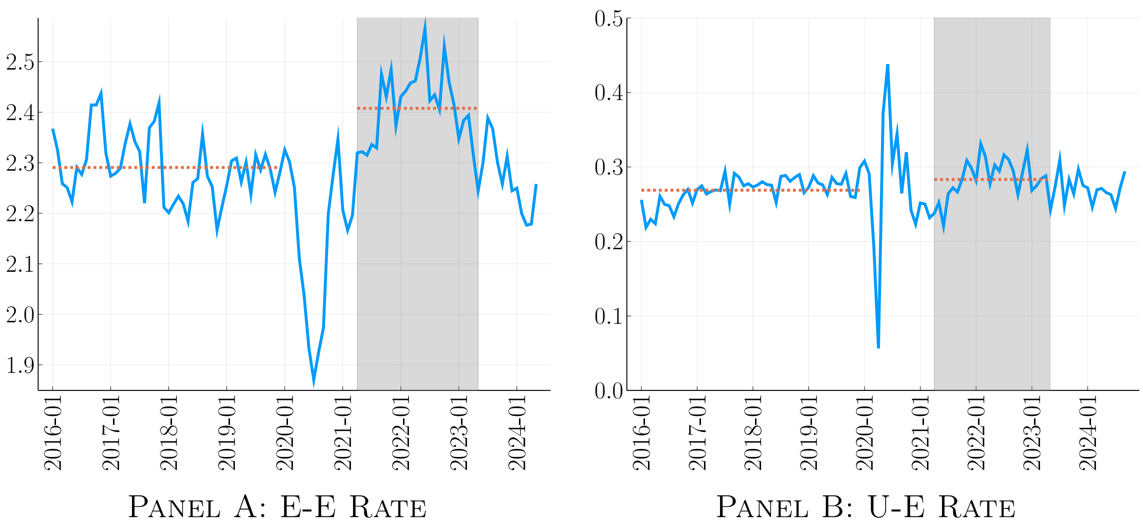

Notes: Panel A of Figure shows the time series pattern of monthly E-E flows averaged at the quarterly frequency using data from the CPS. Each observation is a quarter between January 2016 and May 2024. See the Online Appendix for additional details of the data construction. Panel B shows the time series pattern of monthly U-E flows downloaded directly from the FRED database. The dashed red lines in both panels provide the average flows during the 2016Q1-2019Q4 and the 2021Q2-2023Q2 periods.

Panel B of Figure 3 shows the time series patterns for monthly U-E flows during 2014-2024. The monthly job-finding rate measures the share of unemployed workers who transition to employment during a given month.[13] There was no statistically significant change in the U-E rate between the pre-period and the inflation period. Unemployed workers found

[13]For ease of replication, we downloaded this series directly from the St. Louis Federal Reserves Economic Database (FRED). In particular, we downloaded the series "Labor Force Flows Unemployed to Employed" and "Unemployment Level"; both of these series come from the Current Population Survey and are provided at the monthly level. We divide the former by the latter to make the monthly U-E rate. The FRED database does not include a measure of E-E flows so we made our own series for that from the CPS microdata.

employment in a given month at roughly the same 27% rate during both the inflation period and the pre-period. While it is normally the case that changes in the U-E rate explains the vast majority of unemployment dynamics (Shimer (2012)), changes in the job-finding rate explained essentially none of the unemployment dynamics during the 2021-2024 period.[14]

2.3. Real Wage Growth

To measure trends in real wage growth we use data from the Atlanta Federal Reserve's Wage Tracker Index. The Atlanta Fed Wage Tracker Index uses the panel component of the Current Population Survey (CPS) to make a measure of composition adjusted year-over-year change in the worker's per-hour nominal wage on their main job.[15] Given the Atlanta Fed provides a series on nominal wage growth, we create a series of real wage growth by deflating by the CPI inflation rate over the corresponding period. Finally, for comparability across groups and over time, we create a real wage index that takes the value of 1 in January 2016.

Panel B of Figure 1 in the introduction reports the real wage index for the median worker using the Atlanta Fed data. As discussed in the introduction, median real wage growth increased on average by 1.4% per year during the 2016-2019 period. The real wage growth is driven by both aggregate productivity growth and a life cycle component due to the workers being sampled one year apart. The median worker has a real wage level in 2024 that is similar to their 2019 wage level. However, this does not imply that workers' real wages have recovered during the inflation period. As noted above, real wages should have grown by about 1.4% per year during the 2020-2024 period if the wage growth from 2016-2019 continued. This implies that the median worker still has real wages in June of 2024 that are 4.4% below their predicted real wage. In summary, real wages have plummeted during the inflation period for the median worker and they have yet to catch up to where their wages would have been given trend wage growth.

In our model, we explore how inflation affects search effort and wages of workers with differing levels of productivity. Figure 4 shows the real wage trends for workers in different wage quartiles over the 2016-2024 period using the Atlanta Fed Wage Tracker data; the Atlanta Fed produces wage series for workers whose initial wage is in different quartiles of the

[14] We show the decomposition of changes in the unemployment rate to changes in the U-E rate and changes in the layoff margin in Appendix Figure B.2.

[15] The CPS structure allows for worker wages to be measured exactly one year apart in the survey's outgoing rotation. For salaried workers, the hourly wage is computed as weekly earnings divided by usual weekly hours worked. For hourly workers, the hourly wage is their reported wage per hour. See the Online Appendix for additional details about the Atlanta Fed Wage Tracker Index.

initial wage distribution. Consistent with the findings in Autor, Dube, and McGrew (2024), workers in the bottom quartile had higher real wage growth on average during the 2016-2019 period (2.1% per year) than did workers in higher wage quartiles. However, as seen from Figure 4, all workers have real wages in June 2024 that are substantially below their predicted levels.[16] The real wage declines relative to trend are largest for workers in higher wage quartiles. Specifically, for workers in the bottom, second, third, and top wage quartiles, real wages are -2.6%, -2.1%, -4.7%, and -6.7% below trend as of June 2024, respectively. These results are consistent with the wage compression during the 2021-2024 period documented in Autor, Dube, and McGrew (2024); all groups experienced real wage declines relative to trend during the inflation period but the declines were largest at the top of the wage distribution.

2.4. Wage Growth, Job-Changers vs Job-Stayers

Finally, we use data from *ADP Pay Insights* to examine the relative wage growth of job-changers vs. job stayers during the inflation period. ADP is a payroll processing company that processes payroll for roughly one-fifth of the U.S. labor market. Given the size of the ADP data, ADP can track the components of compensation over time both for workers who remain with the same firm and for workers who transition from one firm to another. As a result, ADP has a much larger sample of job-changers each month than does the CPS.[17] Our analysis with the ADP data spans the 2020 to 2024 period given that ADP Pay Insights only started publishing earnings growth data for broad groups such as job-changers vs job-stayers starting in 2020.

Panel A of Figure 5 shows the median annualized nominal earnings growth (year-over-year) for (i) workers who remained with their same employer during the prior 12 months

[16] Real wages have positive growth during most periods of time. As a result, it is necessary to evaluate real wage growth dynamics relative to trends. However, there is little theoretical guidance on what that trend rate should be. For our baseline results, we take an approach similar to Autor, Dube, and McGrew (2024) and focus on trends in the years prior to 2019. In the appendix, we show results where the trends are defined over the entire 2000-2019 period. The results for the median worker, and workers in the second, third and fourth wage quartiles are essentially unchanged when using a growth rate defined over the longer period. For these workers, real wage growth during the 2016-2019 period was essentially the same as real wage growth over the 2000-2019 period. That is not the case for workers in the bottom wage quartile whose real wage growth during the 2016-2019 period was much higher than over the longer period. Instead, if we assume that real wages for bottom quartile were predicted to grow at the same rate as the median worker during the 2016-2019 period (instead of their higher growth), workers in the bottom quartile of the wage distribution would still be roughly 0.5% below trend as of June 2024.

[17] We downloaded the data from ADP Pay Insights directly from `https://payinsights.adp.com/`. For additional information on how the ADP payroll data can be used to measure changes in compensation over time for the US population, see Grigsby, Hurst, and Yildirmaz (2021).

Figure 4: Real Wage Growth, By Wage Quartile

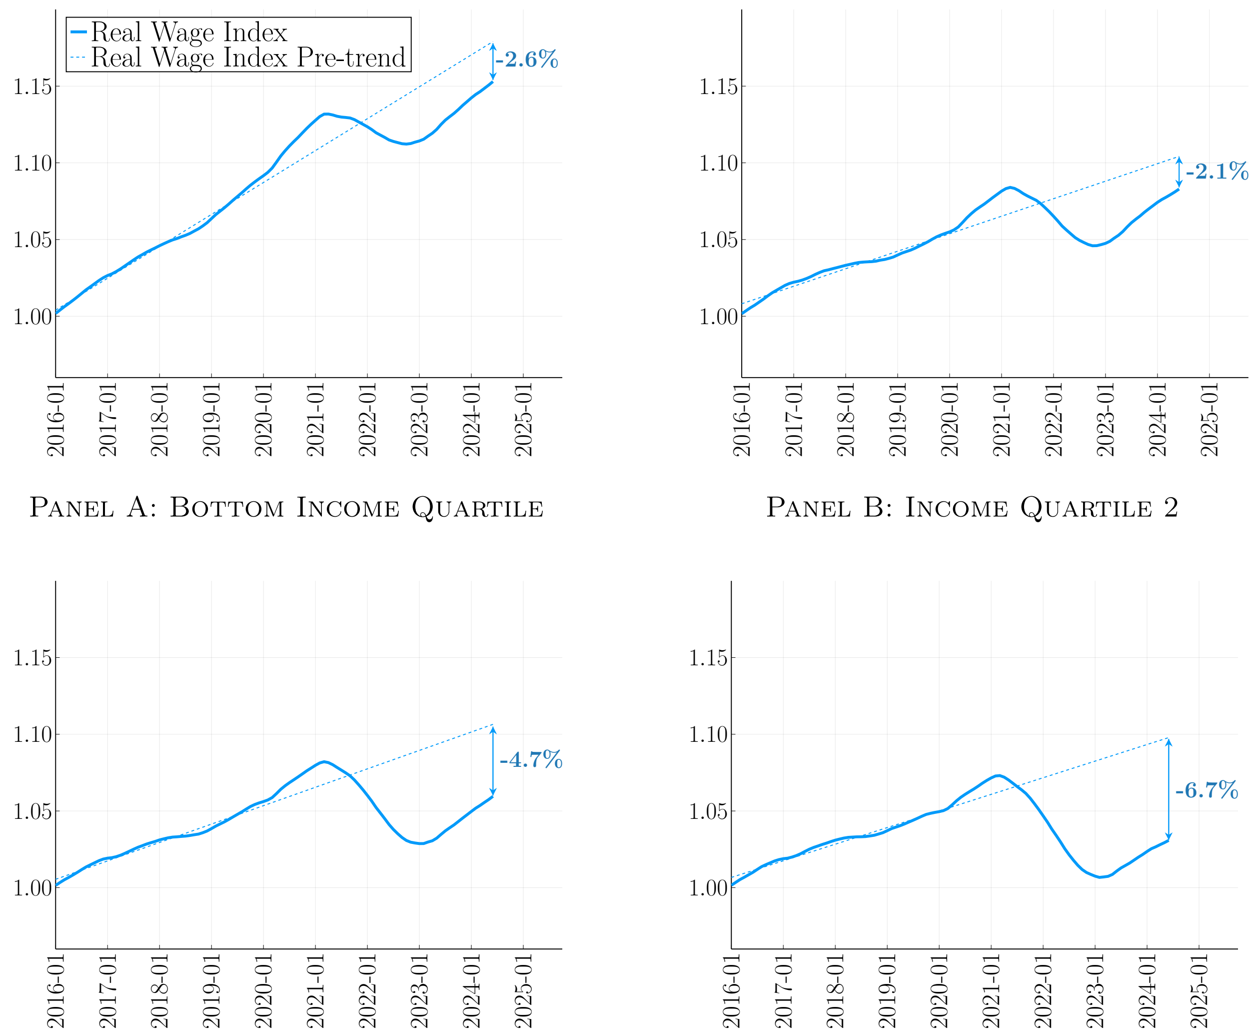

Panel A: Bottom Income Quartile

Panel B: Income Quartile 2

Panel C: Income Quartile 3

Panel D: Top Income Quartile

Notes: The figure shows the evolution of real wages from the Atlanta Fed Wage Tracker Index for workers in different income quartiles. Each figure shows the trend in the real wage based on the 2016-2019 data (in the dashed line). We convert the nominal Atlanta Fed Wage Index into a real wage index by deflating the series for each income quartile by the aggregate CPI.

(job-stayers, dashed line) and (ii) workers who switched employers during the prior 12 months (job-switchers, solid line). During the inflation period, the median nominal earnings growth of job changers increased to over 16%. By mid-2024, the median nominal earnings growth of job-changers appears to have stabilized at around 8.0%. Given that the ADP Pay Insights data started in late 2020, there is no direct way to compare it to a pre-period. However, Grigsby, Hurst, and Yildirmaz (2021) find that median nominal wage growth for job changers

Figure 5: Nominal Wage Growth 2020-2024, Job-Changers and Job-Stayers

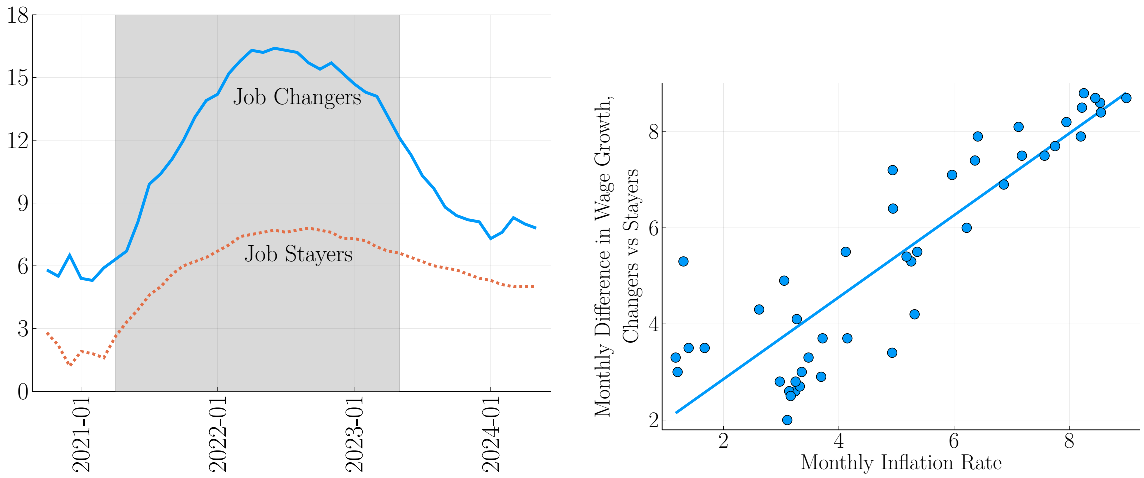

Panel A: ADP Wage Growth, By Switching Status

Panel B: Changer-Stayer Difference vs. Inflation Rate

Notes: Panel A of the figure shows the median nominal income growth of job-stayers (dashed line) and job-changers (solid line) during the October 2020 through June 2024 period from the ADP Pay Insights database. Panel B plots the monthly difference between the two series vs the monthly year-over-year inflation rate. See text for additional details.

in the ADP sample was slightly above 8% during the 2008-2016 period that they analyzed. Given that, the median nominal earnings growth of job-changers in 2024 appears to have returned to pre-pandemic levels. Conversely, the median nominal wage growth of job-stayers peaked at only 8% during the inflation period. Panel B of the figure shows that the gap in wage growth between job-changers and job-stayers is strongly correlated with the monthly inflation rate. As seen from the panel, job changers were able to get even larger wage increases relative to job stayers when inflation was higher.[18]

3 Model

In this section, we develop a model of how workers respond to unexpected changes in the inflation rate and ask whether such changes, all else equal, can causally generate the patterns documented in Section 2. Our goal is to have the model match both the time-series patterns on average and also the heterogeneous patterns across workers of differing types. The model

[18]In the appendix, we show that the Atlanta Fed Wage Tracker data also shows that the gap in the wage growth of job-changers relative to job-stayers roughly doubled during the inflation period.

mixes elements of modern theories of labor market flows with frictions in nominal wage adjustments and lack of commitment on the side of both workers and firms. The interaction of these two frictions leads to infrequent wage adjustments and labor market flows in response to shocks (as discussed in Blanco, Drenik, Moser, and Zaratiegui, 2024). In such an environment, a burst of inflation reduces workers' real wages, as a result of which some workers quit immediately, but those who stay on their job—while becoming more eager to engage in costly renegotiation for higher wages or on-the-job search—become a bargain for their employing firms and are thus less likely to be laid off. Altogether, these incentives result in a burst of quits, higher E-E flows, and higher vacancies for employed workers, consistent with the evidence.

3.1. Environment

Time is continuous and indexed by t. The economy is populated by a unit measure of workers, denoted by $i \in [0,1]$. Workers can either be employed ($E_{it} = 1$) or unemployed ($E_{it} = 0$). Workers die at an exogenous rate $\chi > 0$ and are replaced by newly unemployed workers. To focus on and isolate the effects of rigidities in the labor market, we abstract away from rigidities in firm pricing and assume the price of the homogenous consumption good is exogenous.

Exogenous Worker Shocks. Each worker is subject to an idiosyncratic productivity shock, Z_{it}, that evolve over time as follows. When workers are born, they draw their productivity from a log-normal distribution with mean μ_{z0} and standard deviation σ_{z0} truncated at bounds $\underline{Z} < \overline{Z}$. After birth, worker-specific productivity shock Z_{it} follows a Brownian motion with drift:

$$\mathrm{d}\log(Z_{it}) = \gamma(E_{it})\mathrm{d}t + \sigma \mathrm{d}W_{it}^{Z}, \tag{1}$$

where the drift $\gamma(E)$ potentially depends on the employment state. For example, while employed $\gamma(E)$ may be positive indicating on-the-job human capital accumulation and while unemployed $\gamma(E)$ could be negative indicating a depreciation of skills while not working. The evolution of these productivities is subject to reflecting barriers at the same productivity bounds, $\underline{Z} < \overline{Z}$. We will refer to workers with differing Z's as being workers of differing types.

Production Technology. While employed in a match, worker i produces AZ_{it} units of output where A is an aggregate productivity measure. Such a worker then receives a real wage $W_t = \tilde{W}_{it}/P_t$, where $\tilde{W}_{it}$ is the nominal wage and P_t is the price level with growth rate denoted by π. While unemployed, worker i receives a flow real income of $BZ_{it}^{\phi_B}$, which captures the flow value of non-employment.[19] The parameter ϕ_B measures the extent to which the flow value of non-employment scales with worker productivity. When $\phi_B < 1$, employed low-productivity workers will be, on average, closer to their value of non-employment. Conversely, when $\phi_B > 1$, employed high-productivity workers will be, on average, closer to their value of non-employment. ϕ_B will be one important parameter in determining whether the elasticity of worker flows in response to labor market shocks differs across worker types.

Search and Matching Technology. Job search is frictional and directed on both the worker and firm sides. Firms announce wage-specific vacancies to attract workers with productivity Z at a vacancy posting cost of KZ^{ϕ_K}. There is an infinite mass of potential firms that can open a vacancy and hire a worker at any of these markets. Thus, the expected benefit of opening a vacancy in any market must be zero. The parameter ϕ_K measures the extent to which vacancy posting costs scale with worker productivity. When $\phi_K < 1$, vacancy posting costs will be proportionally smaller for high-productivity workers. Conversely, when $\phi_K > 1$, it is more expensive for firms to hire a high-productivity worker. ϕ_K will be the second important parameter in determining whether the elasticity of worker flows in response to labor market shocks differs across worker types.

The creation of matches in each market is governed by a standard matching function with constant returns to scale between vacancies and the search effort of workers. Each worker chooses search effort s subject to a convex utility cost function that depends on their search effort and employment status, denoted by $S(s; Z, E) = \eta(E)^{1/\phi_s} \frac{s^{1+1/\phi_s}}{1+1/\phi_s} Z$, where $\eta(1) > \eta(0)$ and $\phi_s > 0$. In addition to endogenous separations, matches are also subject to exogenous separation shocks at rate $\delta(Z_{it})$ that possibly varies with worker productivity.

Let $\theta(Z, W)$ denote a measure of tightness in its corresponding market; i.e., the ratio

[19] The fact that the value of non-employment is in real terms is consistent with the findings of Chodorow-Reich and Karabarbounis (2016a). That paper finds that the value of non-employment consists of the value of non-work time (measured in units of consumption) and unemployment insurance benefits. The value of non-work time includes the value of leisure or home production. Using a variety of empirical approaches, Chodorow-Reich and Karabarbounis (2016a) find that most of the value of non-employment is due to the value of non-working time (e.g., leisure). The value of leisure is not subject to nominal rigidities so assuming the value of non-employment is in real terms is consistent with their findings.

of vacancies to the total effective units of search intensity of workers with productivity Z looking in the market with a real wage W. In a market with market tightness θ, workers find jobs with probability $sf(\theta)$, while firms find workers with probability $q(\theta) = f(\theta)/\theta$. As is common in the literature, we assume that $f(\theta)$ is increasing, $q(\theta)$ is decreasing, and that $f(0) = 0$ and $\lim_{\theta \downarrow 0} q(\theta) \to \infty$. We assume that firms and workers can only visit one market at a time.

Wage Determination within a Match. Once in a match, workers's wages are subject to three frictions. First, the adjustment of nominal wages is costly to workers—in units of utility—where they can initiate a wage bargaining process with their employer subject to a randomly drawn fixed cost, ψ. At any point in time, with probability $\beta^+ \mathrm{d}t$, the worker can pay a stochastic cost $\psi^+ Z$ in units of output to increase the current wage. Similarly, with probability $\beta^- \mathrm{d}t$, the worker can pay $\psi^- Z$ units of output to start bargaining to negotiate a wage cut. With the remaining probability, bargaining costs are infinitely large. The cumulative distributions for ψ^+ and ψ^- are $\Psi^+(\psi)$ and $\Psi^-(\psi)$ with non-negative support, respectively. Upon bargaining, the new wage is set according to the Nash Bargaining solution, where the worker's bargaining power is denoted by τ and the outside option in case bargaining fails is the dissolution of the match.[20] Second, in addition to costly bargaining, we also model opportunities for "free" wage increases where, with Poisson arrival rates of β^{Π}, workers receive nominal wage increases that are capped above by a "target inflation rate," denoted by π^*. Finally, neither firms nor workers can commit to staying in a match. This, given the worker's wage at any point in time, allows either party to endogenously dissolve the match through unilateral layoffs or quits.

It is worth noting that the nominal rigidities in this model only occur with respect to the wages of workers within a current match. As is common in the literature, we assume that the wages of new hires are perfectly flexible. This implies that workers can escape their falling real wages on the job when there is a burst of inflation by engaging in costly search for a new match.[21]

[20] It should be noted that we treat search costs and renegotiation costs as being two distinct decisions. However, in reality, they are likely linked. Workers could search for another job and bring their external offer back to their original firm to facilitate a renegotiation of their current wage. In this case, the renegotiation costs could stem, in part, from costly search. Treating search and renegotiation as two separate decisions facilitates model tractability without changing any of the model's broad conclusions. However, given the potential link between search costs and renegotiation costs, we group these costs together when assessing the potential welfare costs of inflation on worker well-being.

[21] There is a recent literature empirically examining the flexibility of new hire wages. Kudlyak (2014) and

Preferences and Payoffs. Workers born at time t have the following preferences over consumption C_{is}, on-the-job search costs S_{is}, and renegotiation costs $R_{is} \equiv \mathbb{I}_{E_{is}=1}\mathbb{I}_{\tilde{W}_{is}\neq\tilde{W}_{is^-}}\psi_{is}Z_{is}$, and they discount the future at rate ρ:

$$\mathbb{E}_t\left[\int_t^{\infty} e^{-(\rho+\chi)(s-t)}\left(C_{is}-S_{is}-R_{is}\right)\mathrm{d}s\right]. \tag{2}$$

On the firm side, the value of being in a match with a worker at time t is given by:

$$J_t(Z_{it},W_{it}) = \max_{\mathcal{T}}\mathbb{E}_t\left[\int_t^{t+\mathcal{T}} e^{-(\rho+\chi)(s-t)}\left(Z_{is}-W_{is}\right)\mathrm{d}s\right], \tag{3}$$

where $\mathcal{T}=\min\{\mathcal{T}_j,\mathcal{T}_h,\mathcal{T}_\delta\}$ is a stopping time describing the match duration before the firm, the worker, or nature dissolves it, respectively. Moreover, a vacant firm's value in submarket (Z,W) is given by:

$$V_t(Z,W) = -KZ^{\phi_K} + q(\theta(Z,W))J_t(Z,W). \tag{4}$$

Equilibrium Definition. An equilibrium for this economy is a set of policy functions for all firms and workers, such that, (i) Given a firm's layoff policy, the workers' policies during employment and unemployment spells are optimal; (ii) given an employed workers' quit, wage renegotiation, and on-the-job search policies, the firms' layoff strategies are optimal, and (iii) the free entry condition for vacancy posting holds in all open submarkets, i.e., $V_t(Z,W)=0$ when $\theta(Z,W)>0$.

3.2. Equilibrium Characterization

In this section, we derive the conditions that characterize the equilibrium of this economy. Let $J(z,w)$, $U(z)$, and $H(z,w)$ denote the values of firms, unemployed workers, and employed workers, respectively, where w denotes the log-real wage and z denotes the log of worker productivity. We now describe the equilibrium conditions for workers and firms.

Bils, Kudlyak, and Lins (2023) provide evidence that new hire wages are more flexible than incumbents. Gertler, Huckfeldt, and Trigari (2020), Grigsby, Hurst, and Yildirmaz (2021), and Hazell and Taska (2024) find that new hire wages seem just as rigid as incumbents. However, these latter papers get most of there identification from recessions when the unemployment rate is high. As with incumbent workers, the flexibility of new hire wages may be asymmetric between periods when nominal wages should fall relative to when nominal wages should increase. Hazell and Taska (2024) actually provides some evidence for such potential asymmetry. We will impose the assumption that new hire wages are more flexible than incumbent wages and see how well this assumption matches the wage changes of job-stayers and job-changers observed during the inflation period. It should be noted that all of the qualitative results in the paper would go through as long as the wages of new hires are more somewhat more flexible than the wages of job-stayers during inflationary periods.

Unemployed Workers. The value of being unemployed is characterized by the following Hamilton-Jacobi-Bellman (HJB) equation for all $z \in (\underline{z}, \overline{z})$:

$$(\rho+\chi)U(z) = Be^{\phi_B z} + \underbrace{\gamma_u \partial_z U(z) + \frac{\sigma^2}{2}\partial_z^2 U(z)}_{\text{Law of motion of } z \text{ during unemployment}}$$

$$\underbrace{+ \max_{s_u, \hat{w}_u} \left\{ s_u f(\theta(z, \hat{w}_u)) \left(H(z, \hat{w}_u) - U(z)\right) - e^z \eta_u^{1/\phi_s} \frac{s_u^{1+1/\phi_s}}{1+1/\phi_s} \right\}}_{\text{Expected value of searching for a job}}, \tag{5}$$

with $\partial_z U(z) = 0$ at the boundaries $z \in \{\underline{z}, \overline{z}\}$ due to reflecting barriers at those points.

The optimal submarket choice $w_u^*(z)$ is the solution to the following problem:

$$w_u^*(z) = \arg\max_{w_u} \left\{ f(\theta(z, w_u)) \left[H(z, w_u) - U(z)\right] \right\}, \tag{6}$$

in which a worker trades off the benefit of finding a job quickly with finding a job that pays a higher wage. The optimal search effort $s_u^*(z)$ is given by:

$$s_u^*(z) = \eta_u^{-1} \left(f(\theta(z, \hat{w}_u^*(z))) \frac{H(z, \hat{w}_u^*(z)) - U(z)}{e^z} \right)^{\phi_s}. \tag{7}$$

Here, η_u^{-1} determines the level of search effort, while ϕ_s is the elasticity of search effort to the expected value of finding a job.

On-the-Job Bargaining. When a worker pays the bargaining cost, the newly renegotiated wage is characterized by the Nash bargaining solution:

$$w_b^*(z) = \max_{w_b} \left(J(z, w_b)\right)^{1-\tau} \left(H(z, w_b) - U(z)\right)^{\tau}. \tag{8}$$

From the optimal bargaining decision, we have the bargaining hazard $\beta(z, w)$ given by:

$$\beta(z, w) = \beta^+ \mathbb{I}_{\{w_b^*(z,w) > w\}} \Psi^+ \left(\frac{H(z, w_b^*(z, w)) - H(z, w)}{e^z} \right) \tag{9}$$

$$+ \beta^- \mathbb{I}_{\{w_b^*(z,w) < w\}} \Psi^- \left(\frac{H(z, w_b^*(z, w)) - H(z, w)}{e^z} \right). \tag{10}$$

Similarly, the new wage resulting from free adjustments, denoted by $w_{\pi^*}^*(w, z)$, maximizes the same bargaining objective but subject to the constraint: $w_{\pi^*} \in [0, 12\pi^*]$.

The Game Between Firms and Employed Workers. We restrict the strategies of a matched firm and worker to be Markovian, seeking a Markov Perfect Equilibrium in the game between the firm and the worker. Once a match is formed, at any point in time, the only payoff relevant variables for the firm and the employed worker are the worker productivity (z) and the real wage of the worker (w). Given these states, the firm's strategy is to choose whether or not to lay off the worker. We denote by $\mathcal{W}^{j*}(z)$ the set of wages where, in a match with productivity z, the firm chooses to continue the match.[22] We let $w_l(z)$ denote the least upper bound of $\mathcal{W}^{j*}(z)$ and refer to it as the layoff threshold.

The strategy of matched worker with productivity z consists of (i) a search intensity $s_e^*(z,w)$ and a submarket $w_e^*(z,w)$ for on-the-job search, (ii) when to pay the bargaining cost, and (iii) the set of wages for which they do not quit, $\mathcal{W}^{h*}(z)$. The continuation set for the worker is described by a quitting threshold $w_q(z)$, defined as the greatest lower bound of wages for which a worker of productivity z is willing to continue the match. Given these strategies, we define the *continuation set* of the game at productivity z as the intersection of wages for which the firms and the worker are both willing to continue the match, $\mathcal{W}^{h*}(z) \cap \mathcal{W}^{j*}(z)$. It follows that for monotonic strategies where $\mathcal{W}^{j*}(z)$ and $\mathcal{W}^{h*}(z)$ are both half-intervals with $w_q(z) < w_l(z)$, the continuation set at productivity z is the interval $(w_q(z), w_l(z))$.

Within the continuation region of the game, an employed worker's value satisfies the Hamilton-Jacobi-Bellman equation:

$$\begin{aligned}
\rho H(z,w) = e^w &+ \underbrace{\partial_z H(z,w)\gamma_e + \frac{\sigma^2}{2}\partial_z^2 H(z,w) - \partial_w H(z,w)\pi^*}_{\text{Law of motion of } (z,w) \text{ during employment}} \\
&- \underbrace{\delta(H(z,w) - U(z)) - \chi H(z,w)}_{\text{Separation and death shocks}} + \underbrace{\beta^\pi \left(H(z, w^*_{\pi^*}(w,z)) - H(z,w)\right)}_{\text{Value of free wage adjustment}} \\
&+ \underbrace{\beta^+ \mathbb{I}_{\{w_b^*(z,w) > w\}} \int \max\left\{H(z, w_b^*(z,w)) - H(z,w) - \psi e^z, 0\right\} \Psi^+(\mathrm{d}\psi)}_{\text{Net value of costly upward wage adjustment}} \\
&+ \underbrace{\beta^- \mathbb{I}_{\{w_b^*(z,w) \leq w\}} \int \max\left\{H(z, w_b^*(z,w)) - H(z,w) - \psi e^z, 0\right\} \Psi^-(\mathrm{d}\psi)}_{\text{Net value of costly downward wage adjustment}} \\
&+ \underbrace{\max_{s_e, w_{jj}} \left\{ s_e f(\theta(z, w_{jj})) \left(H(z, w_{jj}) - H(z,w)\right) - e^z \eta_e^{1/\phi_s} \frac{s_e^{1+1/\phi_s}}{1+1/\phi_s} \right\}}_{\text{Expected net value of on-the-job search}},
\end{aligned} \tag{11}$$

and for all states where either agent decides to terminate the match, the employed worker's value equals the unemployment value $H(z,w) = U(z)$. Additionally, the standard value matching condition holds at the indifference point of both agents $H(z, w_l(z)) = U(z)$ and

[22] As in Blanco, Drenik, Moser, and Zaratiegui (2024), we require the continuation set to be a weakly dominating strategy to ensure the uniqueness of equilibrium.

$H(z, w_q(z)) = U(z)$. Finally, since the worker chooses the quitting threshold optimally, the smooth pasting condition holds at this point for both state variables, $\partial_z H(z, w_l(z)) = \partial_z U(z)$ and $\partial_w H(z, w_l(z)) = 0$.

Similarly, the HJB equation for a firm employing a worker at wage w with productivity z in the continuation set of the game is:

$$\begin{aligned}\rho J(z,w) = {} & e^z - e^w + \partial_z J(z,w)\gamma_e + \frac{\sigma^2}{2}\partial_z^2 J(z,w) - \partial_w J(z,w)\pi^* \\ & + \beta(z,w)\left(J(w_b^*(z,w),z) - J(z,w)\right) + \beta^\pi\left(J(z,w_{\pi^*}^*(z,w)) - J(z,w)\right) \\ & - \left(\delta + \chi + s_e(z,w_{jj}^*(z,w))f(\theta(z,w_{jj}^*(z,w)))\right).\end{aligned} \tag{12}$$

For $w < w_q(z)$ or $w > w_l(z)$, we have that $J(z,w) = 0$. The value matching and smooth pasting conditions are $J(z, w_l(z)) = J(z, w_q(z)) = 0$ and $\partial_z J(z, w_l(z)) = \partial_w J(z, w_l(z)) = 0$, respectively.

Finally, the boundary conditions for the firm and worker values at the reflecting barriers are given by $\partial_z J(z,w) = \partial_z H(z,w) = 0$ for $(z,w) \in \{\underline{z}, \overline{z}\} \times \mathbb{R}$.

4 Quantifying the Model

In this section, we calibrate the model parameters and discuss the behavior of workers and firms in the steady state of the model under these calibrated values.

4.1. Calibration

We calibrate the model using the simulated method of moments (SMM) approach and by matching several moments of the microdata. Table 2 shows the assigned values for all parameters of the model under our calibration strategy. Table 3 and Figure 6 demonstrate the goodness of fit between the model and our targeted moments. The time period in our model is one month.

Preferences. We set the monthly discount factor ρ to 0.005, consistent with an annual discount rate of 6% (Hall, 2017). The death rate χ is calibrated to 0.004 to match the 85th percentile of the labor market experience distribution of 40 years (Durante, Larrimore, Park, and Tranfaglia, 2017).

Productivity process. Initial productivity parameters μ_{z0} and σ_{z0} are chosen to normalize the mean log initial productivity to zero and match a P90-P50 weekly earnings ratio for workers aged 25-27 of 2.12 between 2016 and 2019 from the CPS. For productivity dynamics, we set

Table 2: Model Parameters

Parameter	Description	Value	Target
ρ	Discount factor	0.005	Annual discount rate of 6%
χ	Death rate	0.004	85th perc. of experience dist.
μ_{z0}	Mean of initial productivity	0.0	Normalization
σ_{z0}	Std. of initial productivity	0.58	P90-P50 wage ratio for workers aged 25
γ_e	Productivity drift for employed	0.002	Avg. earnings growth between age 25 to 55
γ_u	Productivity drift for unemployed	-0.006	Earnings loss per month of unemployment
σ	Std. dev. of productivity shock	0.04	P90-P50 wage ratio for workers aged 25-55
α	Elast. of the matching function	0.5	Standard value
K	Vacancy cost	4.91	Avg. job-finding rate
B	Unemployment income	1.364	Avg. endog. separation rate
δ_0	Exog. separation rate function	0.006	Avg. exog. separation rate by income
δ_1	Exog. separation rate function	0.023	Avg. exog. separation rate by income
δ_2	Exog. separation rate function	-2.11	Avg. exog. separation rate by income
η_e	Search cost scale when employed	9.456	Avg. job-to-job rate
η_u	Search cost scale when unemployed	1.0	Normalization
ϕ_k	Elast. of vacancy cost wrt. z	1.38	Avg. job-to-job rate by income
ϕ_b	Elast. of unemp. income wrt. z	0.85	Avg. job-finding rate by income
ϕ_s	Elast. of search cost	0.096	Elast. of search effort wrt. to wages
β_{Π^*}	Prob. of wage renegotiation due to $\bar{\Pi}$	0.083	Avg. arrival of 1 per year
β_+	Prob. of positive wage renegotiation	0.076	Frequency of positive wage changes
β_-	Prob. of negative wage renegotiation	0.008	Frequency of negative wage changes
λ_+	Prob. of free positive wage renegotiation	0.9	Dist. of on-the-job wage changes
ι_+	Mean cost of positive wage renegotiation	0.4	Dist. of on-the-job wage changes
τ	Worker's bargaining power	0.5	Standard value
$\bar{\pi}$	Annual trend inflation	0.02	Annual inflation rate 2016-2019
π^*	Annual target inflation	0.02	Annual inflation rate 2016-2019

Notes: The table lists the calibrated values of model parameters and their targeted moments.

the productivity drift while employed to $\gamma_e = 0.002$ to capture an average 30-year earnings growth for employed workers of 0.7 (Alves and Violante, 2023). The negative drift for the unemployed $\gamma_u = -0.006$ matches the elasticity of wage changes between consecutive jobs with respect to the length of the intervening unemployment spell of -0.006 as estimated by Jarosch (2023). The standard deviation of productivity shocks, σ, is set to 0.039 to match the P90-P50 weekly earnings ratio for workers aged 25-55.

Labor Market Flows. Our goal is to replicate not only aggregate flows but also average flows along the income distribution. Two parameters play an important role in shaping the heterogeneity of these flows: ϕ_K and ϕ_B, which determine how vacancy posting cost and home production of the unemployed scale with workers' productivity, respectively. The ability of unemployed workers to find jobs is determined by their search effort and the job-finding

Table 3: Comparison of targeted moments between model and data

Moment	Data	Model
Frequency of on-the-job wage decreases	0.004	0.004
Frequency of on-the-job wage increases	0.063	0.064
Share $\Delta w_b \in (0,6)/(0,\infty)$	0.73	0.7
Share $\Delta w_b \in [6,11)/(0,\infty)$	0.14	0.15
Share $\Delta w_b \in [11,\infty)/(0,\infty)$	0.13	0.15
Search effort-wage elasticity	-0.52	-0.57
P90/P50 real wages (age 25)	2.12	2.09
P90/P50 real wages (ages 25-55)	2.57	2.53
Avg. 30-year wage growth	0.7	0.66
New wage-unemployment length elasticity	-0.006	-0.006

Notes: The table shows the set of moments (excluding the flows, the results for which are reported in Figure 6) that were targeted for calibration.

rate per efficiency unit of search. To reproduce these flows in the model, we first normalize the search cost of the unemployed to $\eta_u = 1$. The elasticity of the matching function α is set to the standard value of 0.5 from Petrongolo and Pissarides (2001). Then, we choose the elasticity of home production with respect to worker productivity ϕ_B to match the job-finding rate (U-E rate) in the 2016-2019 period across deciles of the weekly earnings distribution from the Basic Monthly CPS. Next, we calibrate the search cost of the employed μ_e, the average vacancy positing cost K, and the elasticity of home production with respect to productivity ϕ_K to match the job-to-job transition rate in the aggregate and across the weekly earnings distribution.

Two observations about the values of these parameters are key for the results we describe below. First, the fact that in the calibrated model $\phi_K > 1$ implies that the vacancy cost of hiring more productive workers is higher *relative* to their productivity. Intuitively, this implies that, all else equal, there are fewer vacancies—in relative terms—for more productive workers which helps the model to match the lower job-finding rate of the higher income workers conditional on E-E transitions in Panel B of Figure 6. Second, the value of $\phi_B < 1$ implies that home production of the unemployed scales less than one to one with their productivity. This means that more productive workers lose more—in relative terms—by staying in the unemployment state and all else equal and searching more intensely relative to others when unemployed. This higher search intensity among the more productive unemployed workers offsets the effect of fewer vacancies posted for higher productivity workers and helps the model to match the higher job-finding rate of the higher income workers conditional on U-E transitions in Panel A of Figure 6.

Figure 6: Targeted Moments: Flows in the Labor Market

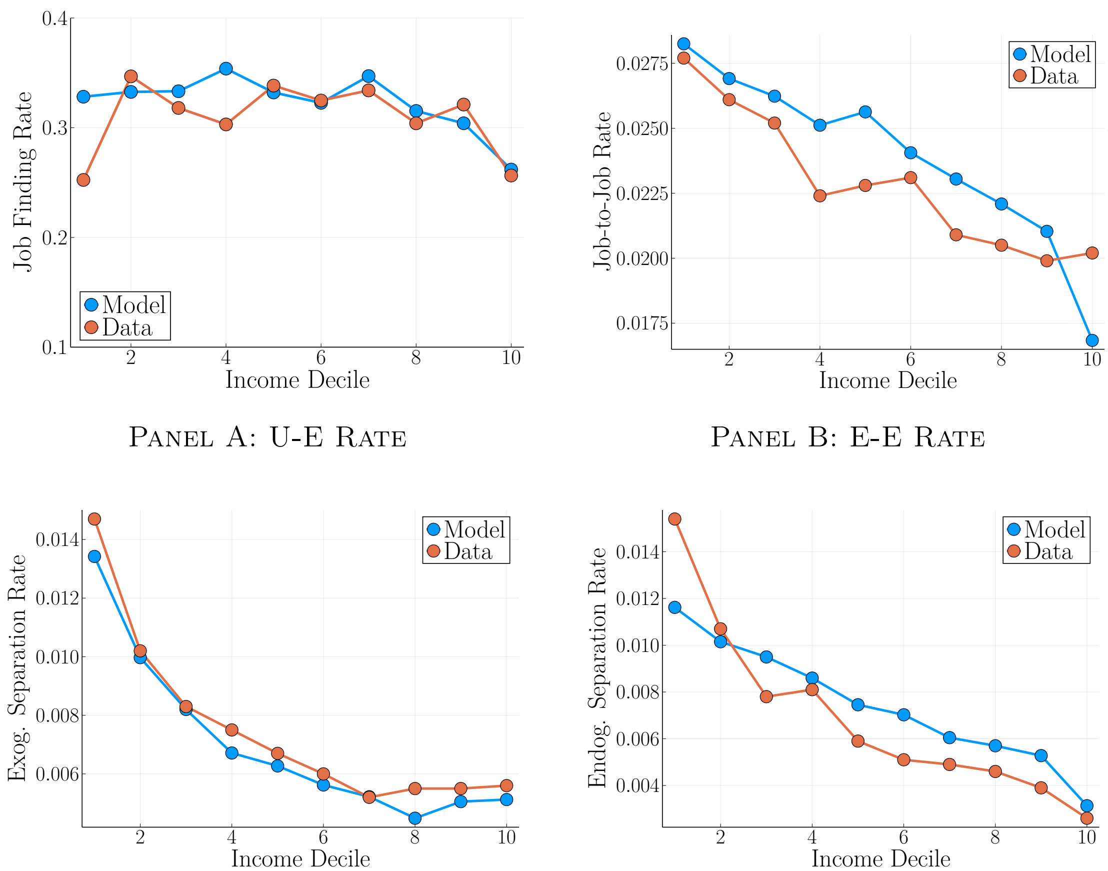

Panel A: U-E Rate

Panel B: E-E Rate

Panel C: Exogenous E-U Rate

Panel D: Endogenous E-U Rate

Notes: The figure shows the flows from employment to unemployment both in the data and as predicted by the calibrated model. These were targeted moments in the calibration. Panels A and B show the job-finding rates of the unemployed and the employed workers, respectively (flows into employment). Panels C and D show the endogenous (quits and layoff) and exogenous separation rates, respectively (flows into unemployment).

In the model, separation into unemployment results from endogenous choices and exogenous shocks $\delta(Z)$. In the CPS, unemployed respondents are asked the reason why they became unemployed, and possible answers are due to a “quit”, “layoff”, or “other reasons”. We map these three motives in the model by considering quits and layoffs as being the outcomes of agents’ choices and interpreting separations due to other reasons as exogenous. We parameterize the relationship between exogenous separations and worker productivity with the following functional form $\delta(Z) = \delta_0 + \delta_1 \exp(\delta_2 Z)$. To set these three parameters, we target the separation rate into unemployment due to "other reasons" across the weekly earnings distribution (see Panel C of Figure 6). To match the average endogenous separation rate given by the sum of quits and layoffs, we exploit the fact that a larger level of home production B raises the opportunity cost of employment, and pushes up the wages that workers search for during unemployment, which gets them closer to the layoff threshold. The calibrated value of B, together with an elasticity ϕ_B, implies a ratio of average home production among the unemployed to average production among the employed of 73%, which is close to the midpoint reported by Chodorow-Reich and Karabarbounis (2016b).

On-the-job Wage Adjustments. We parameterize how wages adjust on the job using the moments provided in Grigsby, Hurst, and Yildirmaz (2021) who use data from the payroll processing firm ADP to measure wage adjustments for U.S. workers during the 2008-2016 period. In particular, the paper shows the following: (i) essentially no job-stayer gets a nominal wage cut during a year, (ii) about one-third of job-stayers get no nominal wage change during a given year, (iii) about 10% workers get annual wage changes between 0 and 2%, (iv) about one-third of workers get an annual wage change of about 2 or 3%, and (v) there is a long tail of larger wage changes with a drop-off after 3%. This distribution is what motivates us to include both a Calvo and menu cost component to wage adjustments with an asymmetry between wage increases and wage cuts. The Calvo parameter β^{Π} governs the arrival rate of costless wage changes between 0 and π^*. This process helps us match the large spike in wage changes at 0 and 2-3% with a missing mass in between. The heterogeneous menu cost part of the model helps us match the long tail in wage changes of job-changers above 3%.

Specifically, we parametrize the distribution of renegotiation costs for wage increases (Ψ^+) as an exponential distribution with rate parameter ι_+, occurring with probability $1 - \lambda_+$, and zero with probability λ_+. For wage cuts (Ψ^-), renegotiation costs are always zero at a rate of β_-. The parameters (β_{Π^*}, β_+, β_-, λ_+,λ_-, ι_+) mainly affect the outcomes of the on-the-job wage change distribution. The parameters β_+ and β_- directly inform the frequency of positive and negative wage changes, respectively. The former is calibrated to 0.076, implying that opportunities to negotiate wages upward arrive every 13 months on average consistent with the moments provided in Grigsby, Hurst, and Yildirmaz (2021). The value of the latter is much lower; implying that opportunities to bargain wage cuts arrive every 10 years, which is needed to match the observed small share of negative wage changes found in Grigsby, Hurst, and Yildirmaz (2021).

Conditional on a positive wage change opportunity, the parameters β_{Π^*}, β_+, β_-, λ_+ and ι_+ shape the wage change distribution as discussed above. While β_{Π^*} is set to 0.083 to reflect common human resources practices that nominal wages have the opportunity to adjust nominal wages once a year, λ_+ and ι_+ affect the share of small versus large wage changes (e.g., a larger expected menu cost shifts the distribution toward larger changes). In order to satisfy Hosios condition (Hosios, 1990), we set the worker's bargaining power τ to 0.5, which is standard in the literature.

Steady-State Inflation. Finally, we set both the trend inflation $\bar{\Pi}$ and the target inflation π^* to 2% annually, consistent with the observed inflation dynamics during the post-2000 period within the United States.

Untargetted Moments. Furthermore, as a result of these two forces, our model also implies a negative relationship between markdowns and productivity, which is a non-targeted prediction of our model that aligns well with the findings of Chan, Mattana, Salgado, and Xu (2023) using Danish microdata.[23] Figure 7 shows the average value of markdowns—defined as the log difference between real wage and worker productivity—per productivity decile in the equilibrium of our model. Higher productivity workers, on average, experience lower markdowns because, all else equal, they face lower job-finding rates and they are relatively more unproductive at home than in a job.

Finally, although in the calibration process, we target only the average separation rate, the model matches the pattern of endogenous separations to unemployment across the deciles of income distribution relatively well as seen in Panel D of Figure 6.

4.2. Model Mechanisms

In this subsection, we describe the policy functions of the workers and firms that arise in the equilibrium under the calibrated parameters. In doing so, we plot outcomes based on the two main state variables of workers; i.e., either as functions of productivity—since the cost of hiring and workers' incentives to search differ by productivity—or as functions of their markdowns defined by the log-deviation of the real wage from productivity, $\hat{w} \equiv w - z = \ln(W/Z)$. Given

[23] Volpe (2024) also documents that lower productivity workers face smaller wage markdowns using Norwegian data.

Figure 7: Equilibrium Markdowns by Productivity Decile

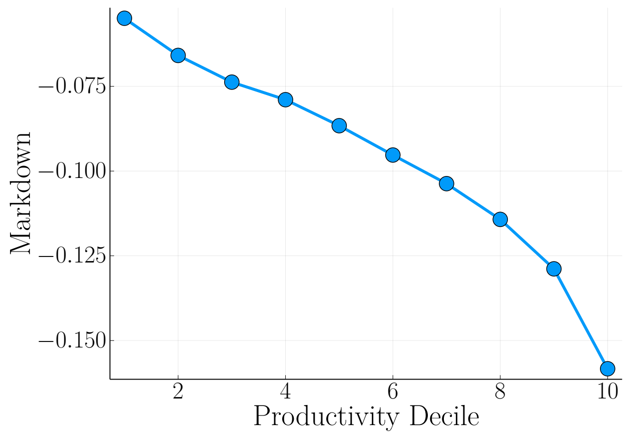

Notes: The figure shows the average markdowns, defined as the log difference between the real wage and worker productivity, for each decile of productivity in the equilibrium.

that we already track the productivity of workers as a state, markdowns are a natural change of variables for wages, as they measure the gap between the wage at which a worker is willing to work and their productivity (which coincides with their marginal product of labor in this model). Moreover, noting that in a competitive model with no frictions, a worker's wage is equal to their marginal product of labor, the deviation of markdowns from zero captures the extent to which wages deviate from such a benchmark.

Job Finding Rates and Starting Wages Conditional on U-E Transitions. We start by discussing the policies of the unemployed. To recap their decisions, at any given point in time each unemployed worker with productivity z targets a particular wage by choosing a submarket and decides how intensely to search for a job. In choosing which submarket to target, these workers internalize that given their productivity, markets with lower markdowns (wages relative to productivity) have higher job-finding rates. This is why a worker would potentially be willing to target a submarket with a lower markdown (i.e., a wage farther below their productivity) if the job-finding rate is high enough.

Panel A of Figure 8 shows the search effort of the unemployed workers as a function of their productivity. The search effort is increasing in productivity because the cost of staying unemployed is higher for more productive workers. This last observation is an artifact of $\phi_B < 1$, which implies that the ratio between the productivity of unemployed workers at home vs. in a match gets larger as z increases.

Panel B of Figure 8 shows markdowns at which the unemployed workers start working

Figure 8: Search Effort and Wages Conditional on U-E Transitions

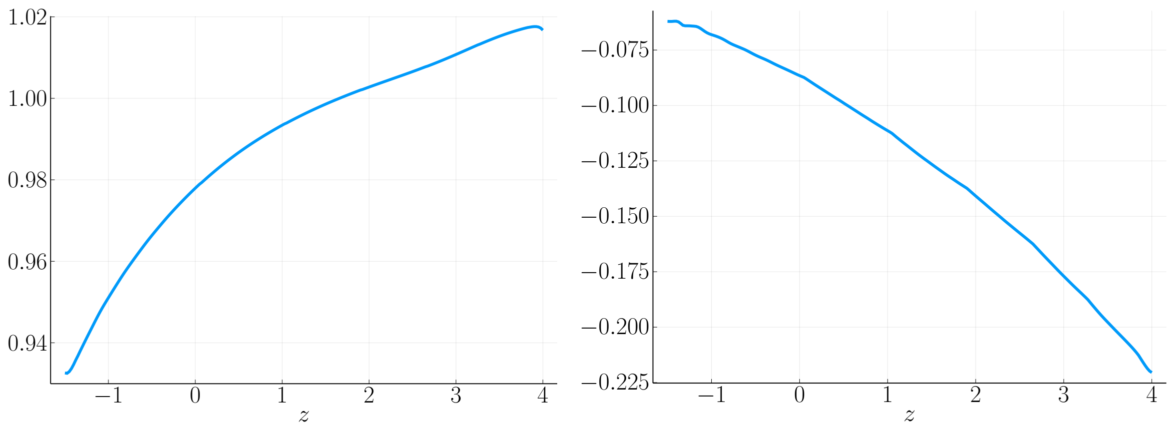

Panel A: Search Effort

Panel B: Starting Markdowns

Notes: Panel A of the figure shows the equilibrium search effort of the unemployed workers as a function of their productivity. Panel B shows the starting markdowns (defined as log real wage minus productivity) of these workers conditional on finding a job.

conditional on finding a job. Despite higher search effort on the part of more productive unemployed workers, the starting markdowns are decreasing in productivity. This is because the market is thinner for more productive workers as the parameter $\phi_K > 1$ implies the cost of hiring more productive workers is higher relative to their productivity. This means that the more productive workers are willing to accept lower starting wages to mitigate the negative effects of these higher hiring costs on their job-finding rate. The result of these opposing forces is a relatively flat job-finding rate across the income deciles for the unemployed, as shown in Panel A of Figure 6.

Value Functions and Margins of Endogenous Separations. Having specified the search behavior of the unemployed, we now turn to the equilibrium strategies of employed workers and matched firms. Panel A of Figure 9 shows the normalized values of a low (high) productivity employed worker in blue (orange) as a function of their log markdown, $\hat{w} = w - z$. To make these comparable, we have subtracted the value of unemployment from the value of the employed and have normalized them by the worker productivity z. Thus, the plotted value shows the value gained above the value of unemployment per unit of productivity, $(H(w, z) - U(z))/e^z$, for different values of markdown, $\hat{w}$, within the match. Similarly, Panel B of the same figure plots the value of the firm matched to each of these workers per unit of worker productivity;

i.e., $J(w,z)/e^z$ as a function of the markdown within the match, $\hat{w}$.

Figure 9: Normalized Values of Employed Workers and Matched Firms

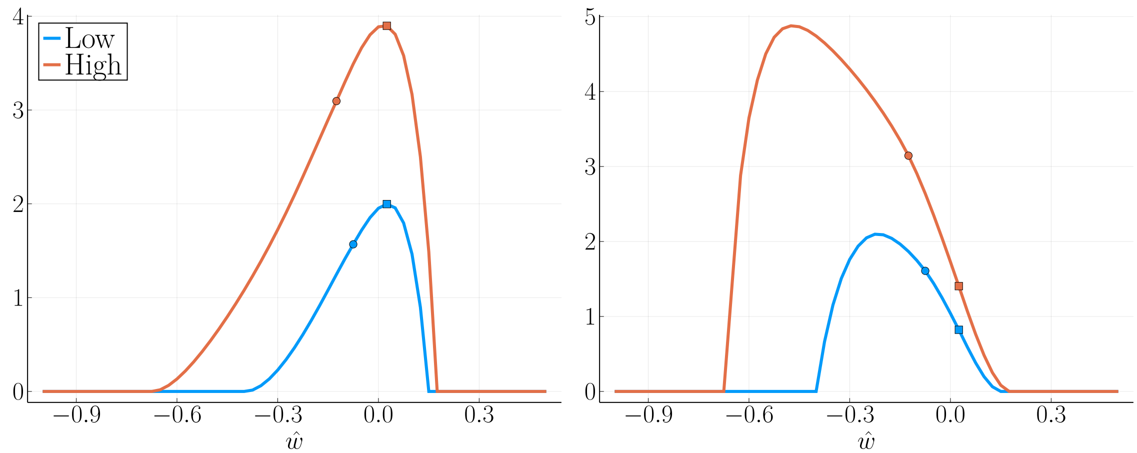

Panel A: Employed Workers　　Panel B: Matched Firms

Notes: Panel A shows the values of a high (low) productivity employed worker in orange (blue) net of their unemployment value and normalized by their productivity, $(H(w,z) - U(z))/z$, as a function of the markdown, $\hat{w} = w - z$. Panel B shows the value of the firm matched to each of these workers per unit of worker productivity, $J(w,z)/z$. In each plot, the circle marks the starting markdown of the match, and the square marks the maximum markdown that the worker seeks within the match.

First, we see that all values are non-monotonic in markdowns. For low enough markdowns, when workers' wages are too far below their productivity, workers would gain from increasing their wages and gaining higher markdowns. It is important to note, however, that these values also encode the future possibilities of being laid off by the employer, should the wage start to far exceed the worker productivity. It is for this reason that, somewhere near the markdown of zero, the value of the worker starts decreasing in markdowns as the worker's wage becomes too close to the wage at which the firm would end the match and lay off the worker. Thus, we would expect workers' values to decline to zero (i.e. $H(.,z) = U(z)$) once for some low enough markdown, where the worker is indifferent between being employed and unemployed, and once at a high enough markdown, where the firm would end the match and the worker would transition to unemployment. These two points, therefore, bound the interval of markdowns in which the match continues. We refer to this interval as the continuation region.

Moreover, to further elaborate on the incentives of the workers, all plots mark the starting markdown of the match—i.e., the markdown at which the match initiates—with a circle, which corresponds exactly to the markdowns at which the unemployed workers start working in Panel B of Figure 8. As discussed above, more productive workers are willing to accept lower starting wages to mitigate the negative effects of higher hiring costs on their job-finding rate. Furthermore, in each plot, we have also marked the highest markdown that the worker would seek within the match with a square. Such a wage is within the interior of the continuation region because of workers' tendency to avoid increasing their wages too far above their productivity, as this would increase the risk of being laid off by the firm.

A final observation about this figure is that the values of the more productive workers lie above those of the less productive workers, even beyond the normalization of the values by productivity. To explain why, it is important to observe that in a model where vacancy and home production are both scaled with productivity—i.e., $\phi_B = \phi_K = 1$—then both plots should exactly be the same (Blanco, Drenik, Moser, and Zaratiegui, 2024), meaning that both workers should earn the same value per unit of productivity. However, under our calibrated model, more productive workers gain more from their match per unit of productivity for two reasons: first, their relative productivity during unemployment is dampened by the fact that $\phi_B < 1$, and second, the fact that $\phi_K > 1$ implies that the chances of finding a new job from unemployment are lower. Put together, these two effects imply that less productive workers are more willing to dissolve the match and quit to unemployment—as seen by the higher lower bound on their continuation region in Figure 10. Such workers are relatively similarly productive when unemployed, and they have a higher chance of finding a new job.

Bargaining and Wage Adjustments on the Job. We now turn to the equilibrium strategies of the workers and firms in the context of on-the-job bargaining. In the model, employed workers have the option to pay a randomly drawn fixed cost to renegotiate their wage with their matched firm. After the menu cost is paid, the new wage is then determined by the Nash bargaining solution, which implies that the wage is a function of the worker's productivity at the time of bargaining.

Therefore, the decision to bargain wages for a given worker with productivity z and markdown $\hat{w}$ depends on the following cost-benefit tradeoff. Such a worker understands that conditional on bargaining, the paid fixed cost will be sunk and can perfectly anticipate their wage change. Thus, the value that is gained by bargaining is independent of the draw of the fixed cost and depends on the differential value under the anticipated bargained markdown and the worker's current markdown. Naturally, this gain in value is smaller if the markdown

Figure 10: Layoff and Quit to Unemployment Thresholds

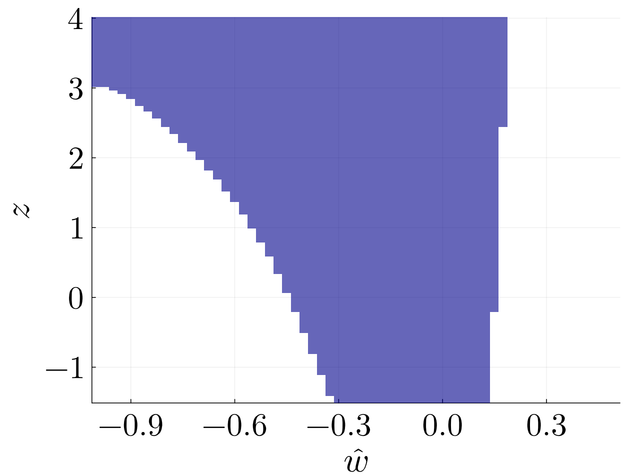

Notes: For every value of z on the y-axis, the corresponding shaded area in the figure shows the continuation set of a match; i.e., the set of markdowns for which both the worker and the firm are willing to stay within the same match. Continuation regions are wider in matches with higher productivity and lower productivity workers have a higher quit to unemployment threshold (the left boundary).

is closer to the worker's optimal markdown within the match. Moreover, since the fixed cost of bargaining, drawn at random by nature, is i.i.d. across workers, only a fraction of workers will find themselves with low enough costs to bargain for these additional values.

Panel A of Figure 11 shows the fraction of workers that bargain as a function of the worker's initial markdown (before bargaining) for two different workers with high (orange) and low (blue) productivity. Panel B shows the corresponding wage changes in percentage points. The bargaining rate decreases for both workers for negative markdowns as the value gained by bargaining decreases with higher markdowns. The sharp drop around the zero markdown is a consequence of the asymmetries in the renegotiation cost distribution for wage increases and decreases. These changes in bargaining rate are indicative of selection effects that are present in the model conditional on bargaining: workers with very low wages relative to their productivity are more likely to bargain for a wage increase.

E-E Rates and Starting Wages Conditional on E-E Transitions. In addition to bargaining, employed workers can also adjust their nominal wages by conducting on-the-job search and moving to a new match. Since search is costly, similar to bargaining, search effort varies across workers with their productivity and increases as their markdown deviates more from their optimal markdowns, which leads to higher E-E rates. Panel C of Figure 11 shows the

Figure 11: Bargaining Rates and Wage Changes Conditional on Bargaining

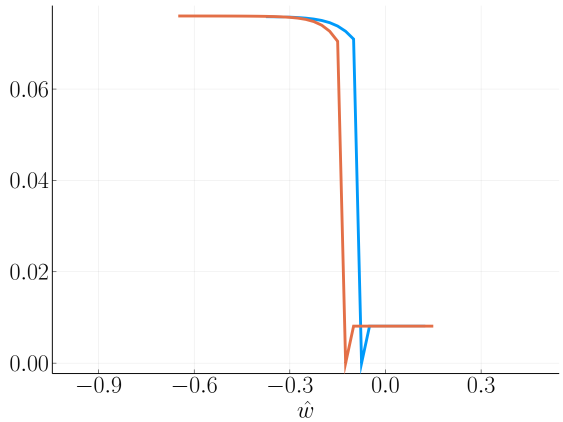

Panel A: Bargaining Rates

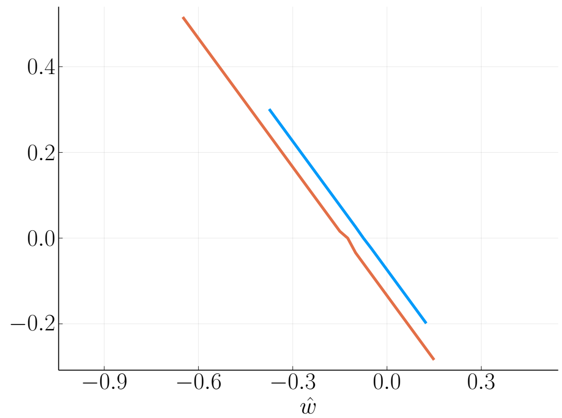

Panel B: Bargained Wage Changes

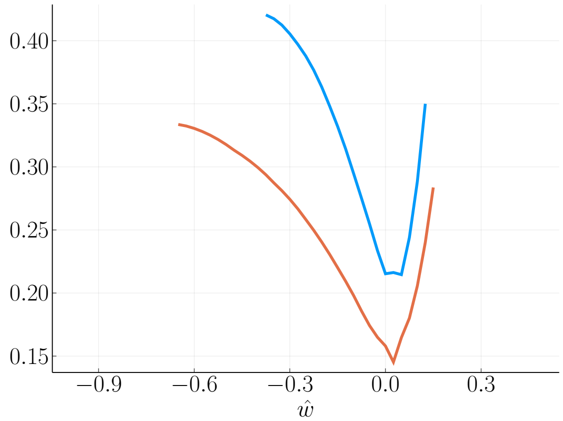

Panel C:
E-E Rate

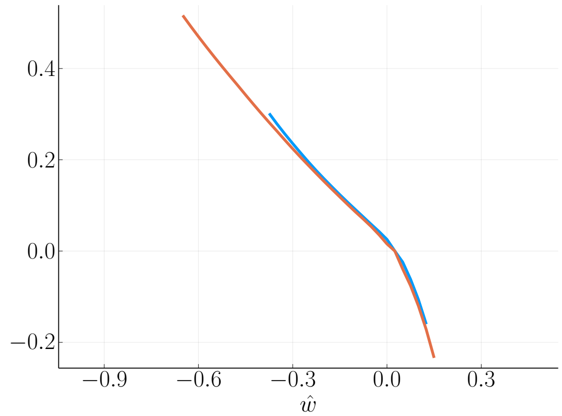

Panel D:
E-E Wage Changes

Notes: Panel A of the figure shows the bargaining rates—i.e. the fraction of workers that bargain for a wage change—of a high (low) productivity worker in orange (blue) as a function of their markdown in the match. Panel B shows the corresponding wage changes (in percentage changes) for these workers conditional on bargaining. Panels C and D show the corresponding E-E rates and wage changes (in percent), respectively.

rate at which a high (low) productivity matched worker transitions to a new job (E-E rate) as a function of their markdown in orange (blue). We see that E-E rate is at its lowest at the optimal markdown and sharply increases as the markdown falls either below or above this optimal level. Panels D of the same figure shows the corresponding wage changes for each productivity type as a function of the worker's current markdown. As described above, low-productivity workers have higher job-finding rates due to the lower vacancy cost of hiring them, which explains why the blue curve lies above the orange curve. Moreover, we see that for each productivity type, their job finding rate increases with the size of their markdown *gap*, which shows the selection effects that are present among job seekers in our model conditional on E-E transitions: The set of workers that seek new jobs is not random and is mostly represented by workers whose wages deviate further from their optimal markdowns.

5 How Workers Respond to Temporary Changes in Inflation

In this section, we assess how labor market flows, the vacancy-to-unemployment rate, wages, and worker welfare respond to a temporary shock to the inflation rate. We start by exploring a one-time unexpected increase in the price level of 13.5%, all else equal. The 13.5% increase represents roughly the jump in the U.S. price level during the April 2021 and May 2023 period. This experiment allows us to assess the dynamics of flows and wages to a one-time shock. We then explore a series of inflation shocks that replicate the time path of inflation we observed during the 2021-2024 period. As seen below, using our calibrated model, the unexpected inflation shock generates patterns for worker flows and wages that match well the actual data highlighted in Section 2 without any additional labor market shocks.

5.1. One-Time Unexpected Increase in the Price-Level

We begin by assessing how a one-time unexpected increase in the price level of 13.5% affects vacancies, quits, layoffs, the vacancy-to-unemployment rate, wage dynamics of job-stayers and job-changers, wage markdowns, and, ultimately, worker welfare. Throughout, we will explore the response of aggregates as well as the effect on workers at different parts of the initial wage distribution.

5.1.1. Wage Markdowns On Impact. Figure 12 shows the distribution of wage markdowns in the economy right before (in green) and right after (in blue) the temporary inflation shock. Given the nominal wage rigidity, an unexpected jump in the price level of 13.5% results in the wage markdown increasing for all workers by 13.5 percentage points upon impact. As a result, the entire distribution of wage markdowns shifts to the left by 13.5% upon impact of the inflation shock. As seen throughout this section, the fact that wages are sticky means that workers become worse off on impact of the shock causing them to take costly actions to have their wages keep up with inflation.

Figure 12: Markdown Distribution Before and After Inflation Shock

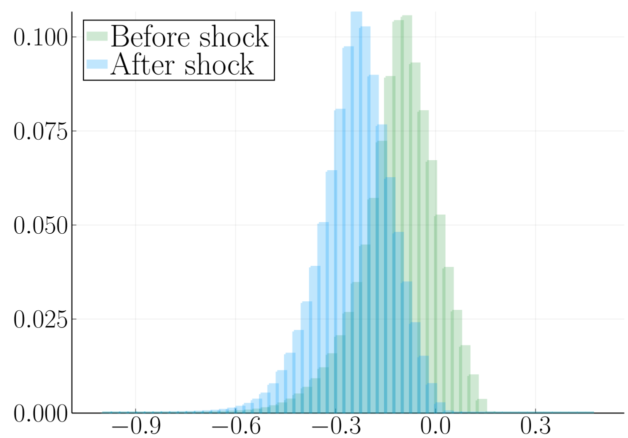

Notes: The figure shows the distribution of wage markdowns right before and right after the unexpected increase in the price level.

5.1.2. Worker Flows, Layoffs, and Unemployment. Panel A of Figure 13 shows how E-E flows respond to the unexpected jump in the price level. On impact, the E-E rate jumps from about 2.1% per month to 3.0% per month; an increase of just under 50%. As shown above, the unexpected temporary increase in inflation causes wage markdowns to increase. In response to this, workers immediately start looking for a jobs at other firms; new hire wages are flexible and will rise in response to the inflation shock. The E-E rate remains elevated for about 18 months in response to a one-time large shock to the price level.[24] Panel B shows how E-E flows respond for individuals in different quartiles of the productivity distribution. Consistent with the data shown in Appendix Table B.1, the E-E rate in the model jumps more for lower productivity workers (first quartile) relative to higher productivity workers (fourth quartile). The reason for this is that lower productivity workers are more elastic in general; they are closer to their outside option and they face lower vacancy posting costs.

[24] On impact, some workers immediately quit into unemployment. These workers were close to the margin between working and not working. These workers then search for a new job from the unemployment pool. This movement will be seen when we show the dynamics of unemployment below.

Additionally, Panel C shows that the U-E rate did not respond to the inflationary shock for any worker type; again this is consistent with the CPS microdata shown in Section 2. Both the returns to working in the home sector and the wages of new hires are set in real terms; a burst of inflation does not alter the relative return between working and not working. To summarize, the model can generate a large increase in E-E flows without a corresponding large increase in U-E flows.

Figure 13: Response of E-E Rate, U-E Rate, and Job-Search to Inflation Shock

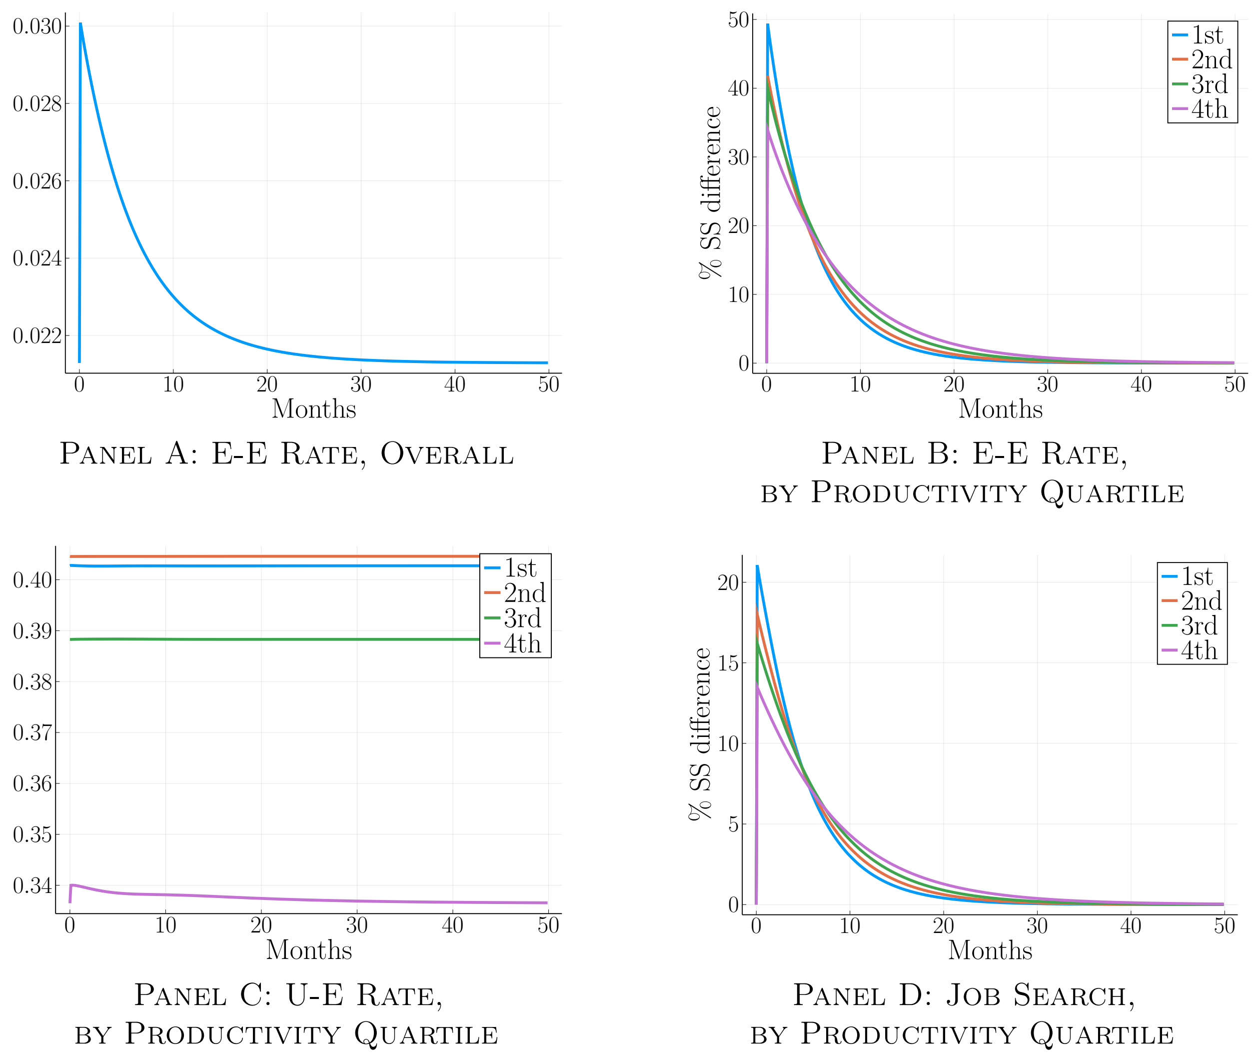

Panel A: E-E Rate, Overall

Panel B: E-E Rate, by Productivity Quartile

Panel C: U-E Rate, by Productivity Quartile

Panel D: Job Search, by Productivity Quartile

Notes: Panel A of the figure shows the time series path of the overall response of the E-E rate to the temporary inflation shock. Panels B to D show the time series response of the E-E rate, the U-E rate, and the job search by quartiles of the productivity distribution.

Panel D shows that the E-E flows come with additional job search costs faced by workers.

In particular, job search increased by 13% on impact for high productivity workers and by 20% on impact for low productivity workers. Notice that for low-productivity workers, the increase in job search is front-loaded. However, for higher productivity workers, their search effort remains elevated for longer periods of time; it takes them longer for their wages to keep up with inflation because they are relatively less elastic. In total, lower and higher productivity workers both substantively increase their search effort in response to a temporary burst in the price level.[25]

Figure 14 shows the response of firm layoffs (Panel A) and the overall unemployment rate (Panel B) in response to the inflation shock. In response to the unexpected shock to the price level, the layoff rate falls sharply. As seen from Figure 12, the markdown distribution shifts to the left after the inflation shocks; workers are now much farther away from the threshold where firms would want to fire a work. The decline in firm layoffs causes the unemployment rate to fall, all else equal. Panel B of Figure 14 shows the time series response of the unemployment rate to the inflation shock. On impact, the unemployment rate rises sharply as a handful of low-productivity workers immediately quit to unemployment in response to the burst of inflation. These workers eventually find new jobs and the unemployment rate declines. However, the unemployment rate continues to decline – dipping below its initial level – as firms lay off fewer workers. The unemployment returns to baseline about two years after the initial shock to the price level.

5.1.3. Vacancy-to-Unemployment Rate and The Beveridge Curve. We motivated the paper by setting out to explain the dramatic increase in the vacancy-to-unemployment rate during the recent period when real wages were falling. As seen from Panel A of Figure 15, a large increase in the price level that causes a temporary increase in the inflation rate of the magnitude that was observed during the 2021-2023 period causally results in an increase in the vacancy-to-unemployment rate. On impact, the vacancy-to-unemployment rate jumps as vacancies dramatically increase from the increased labor market churn; the large increase in vacancies is enough to cause the vacancy-to-unemployment rise despite the unemployment rate increasing. As the unemployment rate falls back towards its initial level, the vacancy-to-unemployment rate increases further peaking at about a 13% increase as vacancies remain elevated. With the large price level shock, the vacancy-to-unemployment

[25] The model prediction that worker job search effort increased in response to the recent inflation is consistent with survey evidence found in Pilossoph and Ryngaert (2023) and Stantcheva (2024).

Figure 14: Response of Layoff Rate and Unemployment Rate to Inflation Shock

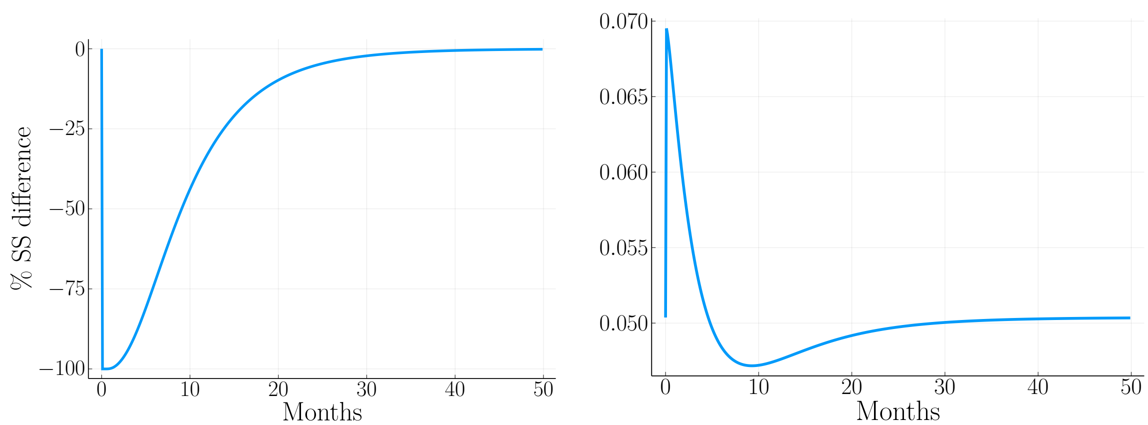

Panel A: Layoff Rate

Panel B: Unemployment Rate

Notes: Panels A and B of the figure show the time series response of the percentage change in the layoff rate and the unemployment rate (relative to the steady-state) in response to the temporary inflation shock.

rate remains substantively elevated for about two years.

Figure 15: The Vacancy-to-Unemployment Rate and Beveridge Curve

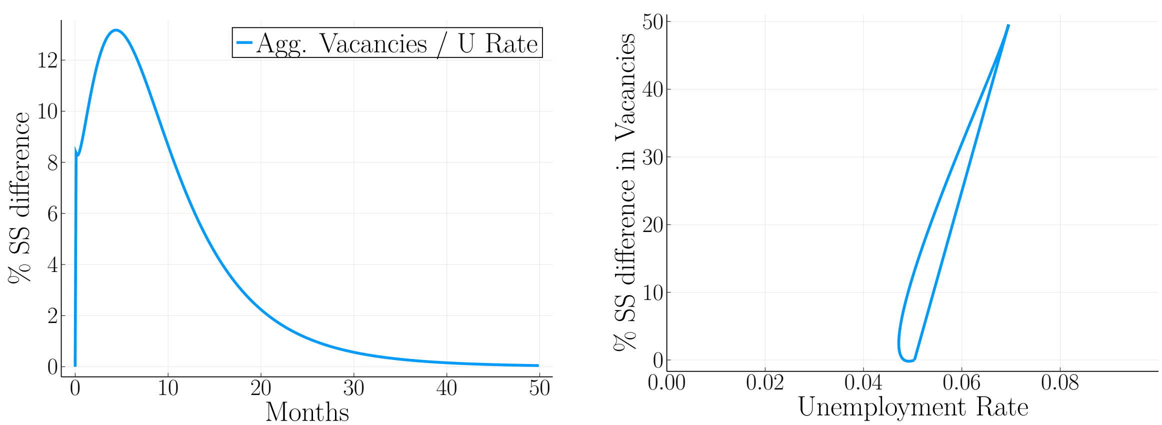

Panel A: Agg. V/U Ratio

Panel B: Beveridge Curve Dynamics

Notes: Panel A shows the time series response of the vacancy-to-unemployment rate in response to the temporary increase in the inflation rate. Panel B shows the dynamics of the Beveridge Curve in response to the unexpected price level increase.

As seen in Panel A of Figure 15, the aggregate vacancy rate (in blue) jumped by about 12% after the inflation shock. The Beveridge Curve plots the relationship between the monthly unemployment rate and the monthly vacancy rate. During the recent inflation period, the Beveridge curve in the United States shifted up and has dramatically steepened. Our model shows how a temporary burst of inflation can explain the recent Beveridge Curve dynamics. Panel B of the figure shows the dynamics of the Beveridge cure in response to the inflation shock. As the price level unexpectedly jumps, aggregate vacancies increase by about 50%. The aggregate vacancy rate is a combination of vacancies filled by workers from unemployment (U-E vacancy rate) and vacancies filled by workers at other firms (E-E vacancy rate). The aggregate rise in vacancies is primarily driven by the increase in vacancies filled by E-E flows.[26] The increase in churn in the labor market as employed workers search for other jobs to increase their real wages results in firms posting more vacancies. Over the next few years, vacancies fall from their peak but are still higher than in the steady state. At this time, the Beveridge Curve can appear upward-sloping as opposed to its usual downward-sloping shape. Unexpected inflation shocks can shift the Beveridge Curve upwards and steepen it because it generates a burst of E-E vacancies without much change in the unemployment rate. In summary, our model can quantitatively explain the upward shift in the Beveridge Curve within the United States during the 2021-2024 period.[27]

5.1.4. Wages. Panels A and B of Figure 16 show the time series response of the wages of job-stayers and job-changers, respectively, to the unexpected increase in the price level. On impact, the wage growth of job-stayers increased by 3 percentage points (from 5% to 8%). For job-changers, wage growth increases by almost 9 percentage points on impact (from 10% to 19%). Again, these patterns are nearly identical to the magnitudes of wage changes of job-stayers and job-changers found in the ADP micro-data highlighted in Figure 5.

Why are the wage changes of job-changers increasing more than the wage changes of

[26] In the appendix, we follow the methodology in Davis, Faberman, and Haltiwanger (2013) to show that the duration of vacancies increased sharply during the inflation period. Our model also generates this finding. In our model, employed workers take longer to fill a vacancy relative to unemployed workers. During the inflation period, there were more employed workers searching for a job and, as a result, the aggregate duration of a vacancy increased.

[27] The shifting Beveridge curve also highlights the difference between this model and benchmark sticky wage models such as the one in Galí (2015), where wage inflation and unemployment are negatively correlated through a conventional Phillips curve. In such a model, a temporary burst of inflation would be mirrored by changes in unemployment, whereas here, a substantial part of the response is through an increase in the aggregate vacancy-to-unemployment rate rather than a change in unemployment.

Figure 16: Wage Change For Job-Stayers and Job-Changers

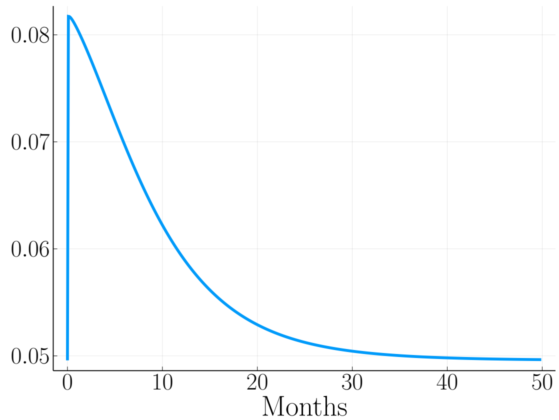

Panel A: Job-Stayers

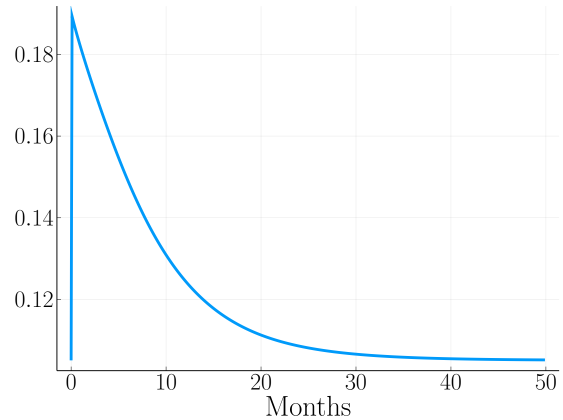

Panel B: Job-Changers

Notes: Figure shows the time series response to the temporary inflation shock on wage growth of job-stayers (Panel A) and the wage growth of job-changers (Panel B).

job-stayers? Job changers are able to have their wage change keep up with inflation because the wages of new hires are assumed to be flexible. For job-stayers, workers have to incur a costly renegotiation with their existing firm. Firms are willing to increase the wages of most of their workers by the steady-state trends inflation rate on average once per year. However, workers can pay an additional cost to try to renegotiate with their firm to get an additional wage increase. Panel A of Figure 17 shows that the model predicts an increase in the fraction of monthly wage changes that occurs after the inflation shock. Data from the Atlanta Fed's Wage Tracker provides evidence that the frequency of wage changes for workers who remain on the job increased sharply during the inflation period. When we compute the full welfare effects of inflation below we need to account for both the additional search costs (which drive E-E wage growth) and the additional renegotiation costs (which drive wage growth for job-stayers) incurred by workers.

Panel B of Figure 17 shows the time series pattern of real income growth in response to the temporary inflation shock for different productivity quintiles. Upon impact, real income falls by 13.5% for all workers.[28] Over time, the real wage of low-productivity workers

[28]For the workers in the lowest decade the real income growth is slightly less than 13.5%. The reason for this is, as noted above, a handful of low-wage workers quit into unemployment after the inflation shock given they were roughly indifferent between working and not working. As these relatively low productivity workers

Figure 17: Frequency of Wage Increases of Job-Stayers and Real Income Response Across Productivity Quartiles

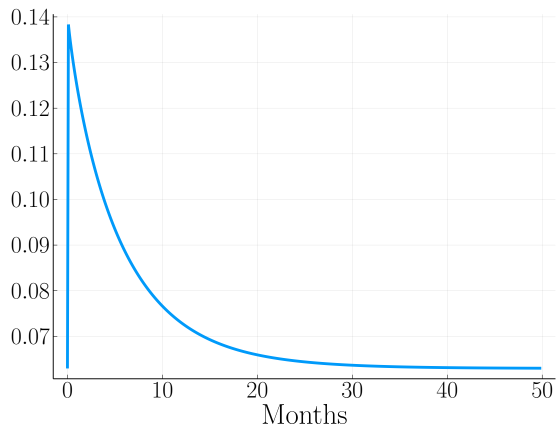

Panel A: Frequency of Monthly Wage Increases of Job-Stayers

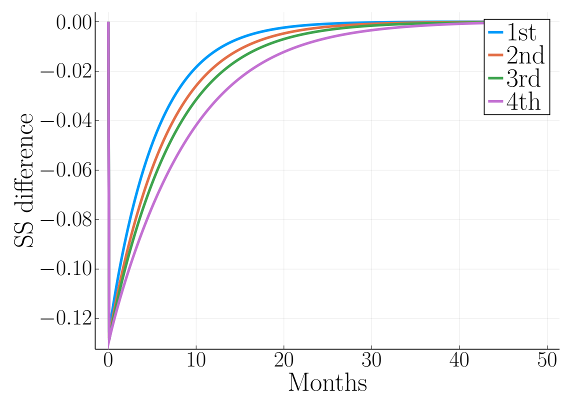

Panel B: Real Income Growth By Productivity Quartile

Notes: Panel A of the figure shows the time series response of the frequency of monthly wage changes for job-stayers in response to the temporary inflation shock. Panel B of the figure shows the time series pattern of the percentage changes in real income in response to a temporary inflation shock for workers in different productivity quartiles. All differences are relative relative to the steady-state path.

recovers faster than higher-productivity workers given that lower-wage workers are more elastic. The higher change in E-E flows of lower-wage workers in response to the inflation shock results in higher real wage growth. Collectively, our model replicates both qualitatively and quantitatively the empirical patterns from the Atlanta Fed Wage Tracker shown in Figure 4. For each productivity quartile, real wages fall but the declines are smallest for lower-productivity workers. As a result, our model generates wage compression during the 2021-2024 inflationary period.

5.1.5. Worker Welfare. Panel A of Figure 18 shows how workers' welfare changes in response to the inflation shock. The welfare costs to employed workers have four components: (i) workers receive real wage declines due to sticky wages in response to the inflation increase (as seen in Figure 17), (ii) workers have to incur search costs to increase their wage at other firms (as seen in Figure 13), (iii) workers have to incur renegotiation costs to increase their

exit the average wage of the remaining workers increases. As a result, the total measured wage loss for this group is slightly less than 13.5% due to a small shift in the composition of the employed.

wage at their own firm (as seen in Figure 17), and workers get a positive benefit given that firm layoffs fall keeping them out of unemployment (as seen in Figure 14). We measure the welfare costs to workers in consumption equivalent units (in multiples of monthly real income before the shock); a welfare cost of 1.0 means a worker would be willing to give up one month of their real wage right before the shock to avoid the temporary increase in inflation. As seen from the figure, workers in each productivity decile are made worse off by the unexpected increase in the price level. However, the welfare loss is larger for higher-productivity workers relative to lower-productivity workers. Our calibrated model estimates that a one-time increase in the price level of 13.5% reduces the welfare of low-wage workers by about eighty percent of a month's income while high-wage workers' welfare is reduced by over one month's income. These results provide a model-based explanation for the survey results in Stantcheva (2024) showing that the vast majority of workers report disliking the current inflationary period.[29]

Panel B of the figure decomposes welfare for workers in the various income deciles into their various sub-components. For all productivity deciles, the primary loss in welfare is due to the real wage declines. This welfare cost reduces the well-being of workers but increases the well-being of those who own the firms. In this case, these welfare losses to workers do not imply an aggregate loss from the inflation. However, the other three components do represent aggregate costs and benefits of inflation. First, notice the orange and green portions of each of the bars. The orange bars represent an aggregate loss from inflation resulting from the costly actions job-stayers have to incur to increase their nominal wage. The green bars represent an aggregate loss from inflation resulting from the costly search undertaken by job changers. The combined welfare losses stemming from costly wage renegotiation and worker search comprise roughly 20 percent of the welfare loss stemming from real wage erosion.[30] However, workers are also made better off as well from the inflation shock. The benefits stem from the fact that workers are less likely to be laid off after the large inflationary shock. In this calibration, these effects are very large because the shock was so big and it immediately moved essentially

[29] Although not shown, we also find that firm values increased sharply during the inflationary period. The loss in welfare to workers from sticky nominal wages provides a net gain to firms. Our results are consistent with the large rise in the profit rate experienced in the U.S. economy during the recent inflation period. As seen in Appendix Figure B.6, the U.S. profit rate increased sharply in the 2021-2023 period; in fact, the U.S. profit to GDP ratio hit a 70-year high in the middle of 2022.

[30] As noted above, we bundle our discussion of search costs and renegotiation costs together given that some of the renegotiation cost can stem from the worker searching for an outside offer to be used in renegotiation with their current firm.

Figure 18: Employed Workers' Welfare Across the Distribution

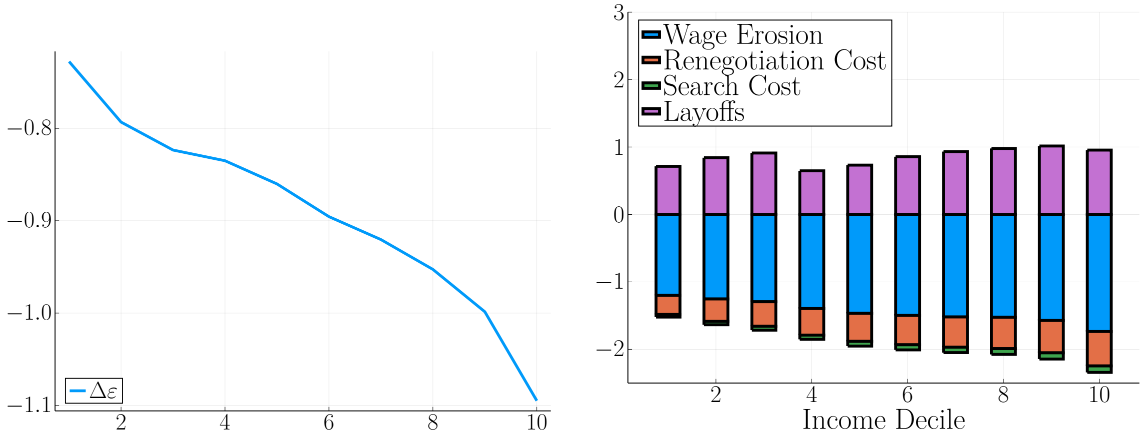

Panel A: Change in Values　　Panel B: Welfare Decomposition

Notes: Panel A of the figure shows the welfare cost of the unexpected inflation shock for workers in different deciles of the worker income distribution. Results are shown in consumption equivalent units of monthly income. Panel B of the figure shows a decomposition of worker welfare into its various components.

all workers far away from the layoff margin. We view this welfare decomposition as being illustrative because the U.S. economy did not experience a one-time increase in the inflation rate of 13.5%; instead, this inflation was spread out over a 26-month period. We turn to that analysis next.

5.2. Series of Unexpected Price Level Shocks

In the prior subsection, we explore the predictions of our model to a large unexpected price level change of 13.5%. That shock was large in the sense that it occurred all at once. However, during the recent inflation period, the price level increased by 13.5% over a 26-month period. In this subsection, we feed in a series of unexpected price level shocks over 26 months that match the price level changes in the U.S. data during the inflation period. When we include a series of price level shocks, we need to take a stance on what the workers and firms are assuming about the future inflation process. For this analysis, we explore the extreme assumption that in each period, the agents expected the trend growth in the price level (i.e., 2% per year) and they were continuously surprised by each subsequent price level shock. The cumulative inflation rate over the 26 months was equal to 13.5%; yet at every point during that period, the individuals expected the price level to increase by by about 0.17% per month (a growth rate of 2% per year).

Figure 19: Response of Labor Market Flows to Series of Inflation Shocks

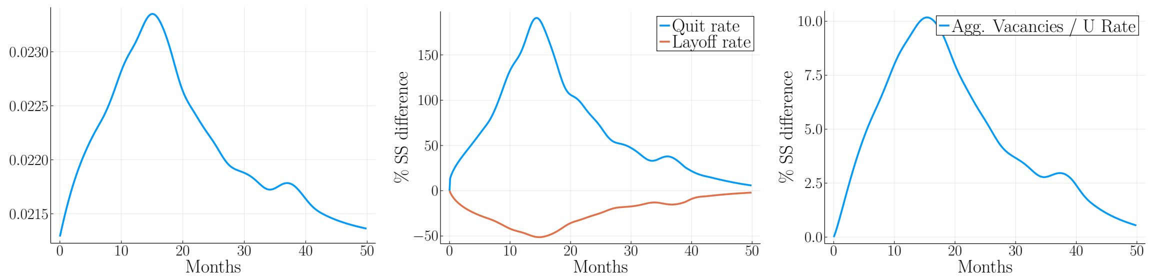

PANEL A: E-E FLOWS

PANEL B: LAYOFFS/QUITS

PANEL C: V/U RATIO

Notes: Figure shows the time series of response to E-E Flows (Panel A), Layoffs and Quit to Unemployment (Panel B), and the Vacancy-to-Unemployment rate (Panel C) from a series of unexpected price level shocks that match the inflation dynamics during the April 2021 to May 2023 period.

Figure 19 shows the response of E-E flows (Panel A), Layoffs and Quits to Unemployment (Panel B), and the Vacancy-to-Unemployment ratio. With the sequence of price level shocks, the flows better match the hump-shaped patterns seen in their data analogs during this period. In particular, all the flows peak about 15 months after the inflation period started. However, the quantitative magnitudes of the change in job-to-job flows and the vacancy-to-unemployment rate are muted in our model relative to the data. Conversely, the layoff rate response and the quits into unemployment response are much larger than we see in the data. The excess layoff and quit margin explain why the E-E flows and the vacancy-to-unemployment response are more muted. As layoffs fall, the unemployment rate actually falls. This puts a downward effect on total vacancies in the economy. As noted above, aggregate vacancies are composed of vacancies attracting workers who are currently unemployed (U-E vacancy) and vacancies attracting employed workers (E-E vacancy). In our steady state calibration, about half of the aggregate vacancies come from U-E vacancies. A sharp decline in the U-E vacancy rate will put downward pressure on aggregate vacancies lowering the V/U ratio.

As seen from panel B, we over-predict the layoff margin in our current calibration. We conjecture that this is occurring because our endogenous separations are too elastic. This comes, in part, because we calibrated our model assuming that about half of the observed separations were endogenous. In the next iteration of the paper, we are going to add some randomness to the separation margins such that our model can also match external evidence on the elasticity of separations to various shocks. We believe this will reduce the layoff margin and, as a result, further increase aggregate vacancies.

6 Inflation and the Vacancy-to-Unemployment Rate Over Time

In the prior section, we documented that during the post-pandemic period, a burst of inflation by itself can cause a sharp increase in the vacancy-to-unemployment rate and an upward shift in the Beveridge Curve similar to the actual magnitudes observed during the 2021-2024 period, all else equal. In this section, we use historical U.S. data to systematically show that periods of high inflation are associated with an increase in vacancies, an increase in the vacancy-to-unemployment rate, and an upward shift in the Beveridge curve after controlling for the unemployment rate.

Figure 20 shows the monthly vacancy-to-unemployment rate in the United States. To make this figure, we use data on aggregate U.S. job vacancies produced in Barnichon (2010). In particular, Barnichon (2010) uses data from the Conference Board's Help Wanted Index and Help Wanted Online Index prior to 2000 and the JOLTS dataset after 2000 to make a harmonized monthly vacancy series for the U.S. between 1951 and 2024. For this figure, we divide the total number of monthly vacancies by the total number of unemployed individuals within the month as reported by the Bureau of Labor Statistics. The average vacancy-to-unemployment rate over the entire 1951-2024 period was roughly 0.7 denoted with the dashed horizontal line in the figure.

As seen from the figure, there are nine periods when the vacancy-to-unemployment rate spiked sharply relative to the average: the early-1950s, the mid-1950s, the late-1960s, the mid-1970s, the late-1970s, the late-1980s, the late-1990s, the late-2010s, and the post-pandemic period. Four of those periods—where the spikes in the vacancy-to-unemployment rate are denoted with the green triangles—are consistent with the traditional view that a rising vacancy-to-unemployment rate represents a "hot" labor market. In particular, the underlying unemployment rate fell sharply as the vacancy-to-unemployment rate rose during each of these periods.[31] These periods were also associated with relatively low inflation rates; the

[31] For example, the unemployment rate fell from about 7% to 4% between both the 1993-1999 period and

Figure 20: Vacancy-to-Unemployment Rate Over Time

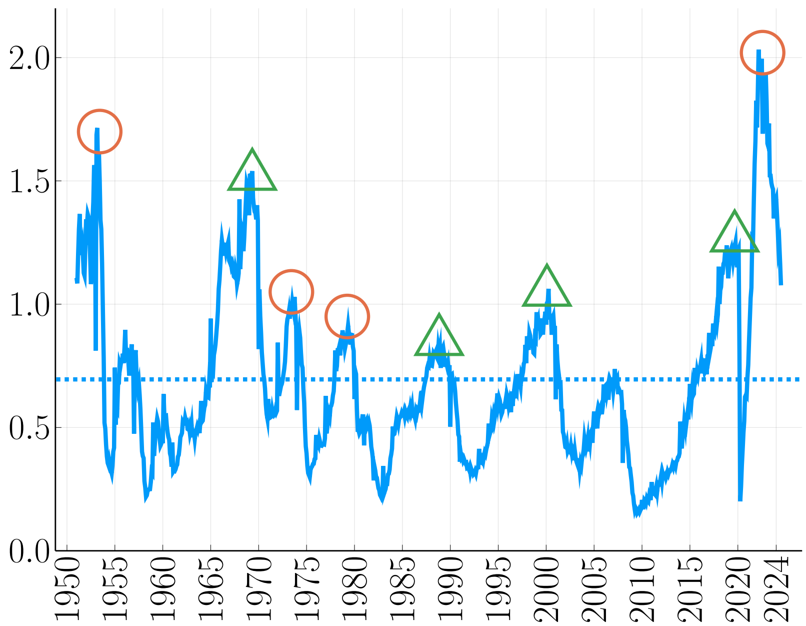

Notes: Figure shows the evolution of the vacancy to unemployment rate over time. There are nine periods since 1950 where the vacancy to unemployment rate exceeded 0.7 for many months. The periods denoted with a triangle are the periods of labor market expansions where inflation was relatively low and unemployment was falling. The periods denoted with the circles are periods when the inflation rate exceeded 7% for multiple consecutive months and the unemployment rate was either high by historical standards (like during the 1970s) or not substantially falling (like the early 1950s and during the recent post-pandemic period).

inflation rate during the green triangle periods was less than 4%. However, four of the other peaks—marked with a red circle— occurred during periods when the inflation rate was persistently above 7% and the unemployment rate was either high by historical standards (in the mid- and late-1970s) or was relatively constant (during the early 1950s and the 2021-2023 period).[32] Notice that the four periods denoted with the red circles are also periods where it has been shown that aggregate supply shocks were important drivers of the observed inflation.[33]

the 2014-2019 period. The unemployment rate fell from about 6% to 3.5% during the 1964-1969 period and from about 7% to 5% during the 1986-1989 period.

[32]The small peak in the mid-1950s had neither falling unemployment nor high inflation so we did not denote it with either a circle or a triangle.

[33]The inflationary period in 1950-1952 was the result of the Korean War, when households scrambled to buy many goods in case there was a return to WWII rationing and supply was constrained given the production of materials needed for the Korean War (see Reed, 2014). The inflation in the mid-1970s was

The above patterns suggest there may be two proximate causes of a rising vacancy-to-unemployment rate. First, the rising vacancy-to-unemployment rate may be caused by a traditional hot labor market story such that labor demand (measured by vacancies) exceeded labor supply (measured by the unemployment rate). During these periods, a primitive positive shock to labor demand puts upward pressure on the vacancy-to-unemployment rate while at the same time putting downward pressure on the unemployment rate; this is the logic underlying the standard downward-sloping Beveridge Curve. However, during other periods, a large burst of inflation may cause excessive labor market churn as workers try to raise their real wages as nominal wages are rigid.[34]

To formally show that high inflation rates can cause a systematic upward shift in the Beveridge Curve and the vacancy-to-unemployment ratio, we estimate the following equation using U.S. monthly data between January 1951 and December 2019 (prior to the start of the global pandemic):

$$y_t = \alpha_0 + \alpha_1 * unemp_t + \alpha_2 * unemp_t^2 + \beta * \pi_t + \epsilon_t, \tag{13}$$

where y_t denotes either the vacancy rate or the vacancy-to-unemployment rate in period t depending on the specification. We define $unemp_t$ as the monthly unemployment rate (in percent) and π_t as the monthly year-over-year inflation rate (in percent). To allow for a potential non-linear Beveridge Curve, we also include the square of the monthly unemployment rate in some specifications. The relationship between the vacancy rate and the unemployment rate is the traditional Beveridge Curve. By including π in the regression, we are assessing whether higher inflation is systematically associated with an upward shift in the Beveridge Curve. We also analyze whether inflation is associated with a rise in measured market tightness by showing how inflation affects the ratio of vacancies to unemployment conditional on the unemployment rate.

The results of estimating these regressions are shown in Table 4. Columns (1)-(3) show the results when the dependent variable is the vacancy rate; these regressions replicate the Beveridge Curve estimates and explore whether inflation systematically shifts the Beveridge Curve. Columns (4)-(6) have the vacancy-to-unemployment rate as the dependent

linked to rising oil prices.

[34] In fact, Hyatt (2015) shows data on job-to-job flows from 1975 through 2013 using data from the Current Population Survey. He documents that job-to-job flows were at their highest level during this 38 year period period during 1979; this was a time when the inflation rate was approaching its highest level in modern U.S. history.

variable, which allows us to explore whether inflation systematically increases the vacancy-to-unemployment rate conditional on unemployment controls. As seen in columns (1) and (4), the unemployment rate itself is a strong predictor of both movements in the vacancy rate and the vacancy-to-unemployment rate. The former is the well-documented Beveridge Curve relationship, while the latter finds that market tightness increases when the unemployment rate is low—the traditional hot labor market story. Columns (2) and (5) highlight the main contribution of our paper. In particular, the results in these columns show that increasing inflation results in an upward shift in the Beveridge Curve by increasing vacancies conditional on unemployment (column (2)) and results in an increase in the vacancy-to-unemployment rate conditional on unemployment (column (5)). For example, the regression shows that an increase in the inflation rate by 10% increases the vacancy-to-unemployment rate by 0.23 percentage points. This is a large effect given that the mean vacancy-to-unemployment rate during this time period is about 0.7. As a reminder, these regressions are estimated using data prior to 2020 suggesting that the link between inflation and vacancies is a common feature of U.S. labor markets during the last 75 years. Finally, columns (3) and (6) show the inflation results persist even when we allow for a non-linear Beveridge Curve by including the square of the unemployment rate in the regression.

Figure 21 shows the partial effect of inflation on the vacancy-to-unemployment rate in graphical form. In particular, we regress the monthly inflation rate on the monthly unemployment rate and the unemployment rate squared and take the residuals from this regression; we call this the residualized inflation rate. We then do the same for the vacancy-to-unemployment rate. The figure plots the two residuals against each other. The slope of the line through the scatter plot is the same as the inflation regression coefficient in column (6) of Table 4. The benefit of the figure is to show that the inflation coefficients in the above table are not being driven by outliers. It is worth noting that the points in the upper-right quadrant of this figure all come from months either during the 1970s or the early 1950s.[35]

Separately, our results are consistent with the relationship between inflation and the vacancy-to-unemployment rate during inflationary periods in other countries. For example, researchers have been studying labor market flows in Argentina around their devaluation in

[35]It should be noted that the corporate profit to GDP ratio was also at historically high levels in 1950, 1974, and 1979. Specifically, between 1950 and 2000, there were only four periods when the corporate profit to GDP ratio exceeded 7%; three of those were the early 1950s, 1974, and 1979. See the online appendix for additional details.

Table 4: Historical Beveridge Curve Estimation

	Dependent Variable: Vacancy Rate			Dependent Variable: Vacancy-Unemp Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Unemp. Rate (%)	-0.251	-0.289	-0.509	-0.152	-0.158	-0.531
	(0.016)	(0.013)	(0.079)	(0.004)	(0.003)	(0.016)
Unemp. Rate Sq.			0.017			0.030
			(0.006)			(0.001)
Inflation Rate (%)		0.142	0.144		0.023	0.026
		(0.008)	(0.008)		(0.002)	(0.001)
R^2	0.24	0.46	0.46	0.67	0.72	0.83

Notes: The Table shows the coefficients from the estimation of equation (13). Each observation is a month between January 1951 and December 2019. Monthly vacancy measure comes from Barnichon (2010) combining data from the Help Wanted Index prior to 2000 and the JOLTS data after 2000. The inflation measure is the year-over-year CPI inflation in a given month. Robust standard errors are in parenthesis. The dependent variable in columns (1)-(3) is the vacancy rate, while the dependent variable in columns (4)-(6) is the vacancy-to-unemployment rate.

2002, which increased the inflation rate to over 30% and increased the unemployment rate to over 20%. Despite the very weak labor market, Albertini, Poirier, and Trupkin (2019) find that the Beveridge Curve steepened in Argentina during the 2002-2004 period. That paper also documents that the vacancy-to-unemployment ratio was increasing during the period of high inflation despite the weakening labor market. Blanco and Drenik (2023) show that the job-to-job rate in Argentina increased as well during this period while real wages were falling. The Argentine experience during the early 2000s provides additional evidence that inflation increases the vacancy-to-unemployment rate by generating an increase in job-to-job flows even when labor markets are not particularly "hot". Collectively, the historical U.S. data and the data from Argentina in 2002 provide additional evidence that inflation itself is associated with an increase in both vacancies and the vacancy-to-unemployment rate, all else equal.

Figure 21: Vacancy-to-Unemployment Residuals vs Inflation Residuals, 1951-2019

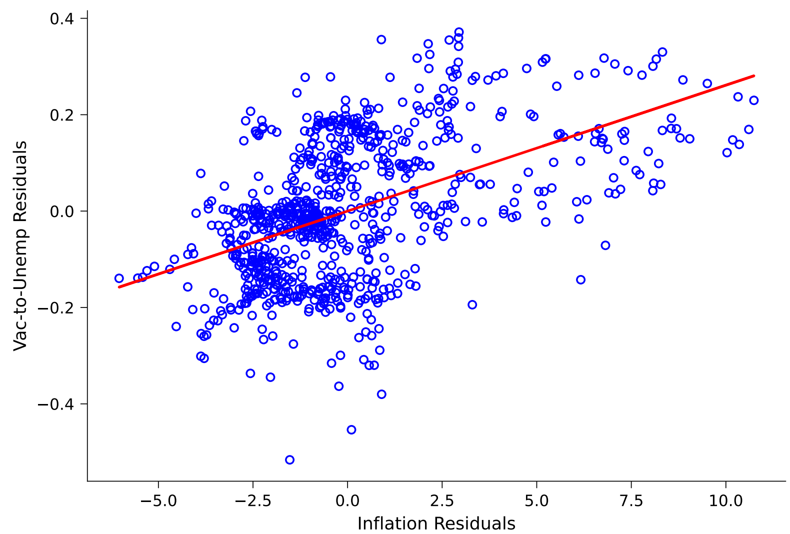

Notes: Figure shows the simple scatter plot of monthly inflation residuals and monthly vacancy-to-unemployment rate residuals. Both the inflation rate and the vacancy-to-unemployment rate are regressed on the unemployment rate and the unemployment rate squared (as in column 2 of Table 4). Each observation is a month during the 1951-2019 period.

7 Alternate Mechanisms for Rising Vacancy-to-Unemployment Rate

Our quantitative analysis above assumed that the only shock to hit the labor market was an unexpected temporary increase in the price level. This allowed us to trace the causal effect of a temporary rise in inflation on labor market flows, wages, and well-being. However, inflation itself is an endogenous variable. In this section, we proceed in two parts. First, we explore how wages and other labor market flows respond to other primitive labor market shocks that can cause a rise in the vacancy-to-unemployment ratio. As we show, traditional hot labor market shocks are unable to jointly match the labor market facts highlighted in Section 2. Second, we specifically discuss the labor market implications of the shocks the literature has identified as the potential cause of the recent inflation.[36]

[36] We thank Ben Moll for encouraging us to add this section.

7.1. Other Labor Market Shocks Through the Lens of Our Model

We begin by exploring a set of other unexpected one-time shocks that can cause the *same* increase in the vacancy-to-unemployment rate as observed in our model. In particular, we define the size of the various other shocks we explore so they roughly match our baseline increase in the vacancy-to-unemployment rate on impact from the one-time 13.5% increase in inflation (shown in Panel A of Figure 15); on impact of our baseline one-time inflation shock, the vacancy-to-unemployment rate increased by 8.5%. Specifically, we explore four different shocks: a positive shock to aggregate productivity (A_t), a negative shock to the household discount rate (ρ), a negative shock to the level of the vacancy posting cost (K), and a negative shock to the value of non-employment (B). Throughout these additional exercises, we still impose the same nominal wage rigidities as in our baseline results. We provide the full details of these exercises in the Online Appendix; here we only summarize the results.

Table 5: Comparison of Alternative Mechanisms That Generate High V-U Rate

Variable	Baseline	Higher Agg. TFP	Lower ρ	Lower K	Lower B
Δ V/U Ratio	8.4	8.4	8.7	8.7	8.4
% Δ EE Rate	41.3	7.5	2.9	5.3	2.6
% Δ UE Rate	0.3	2.9	5.5	4.2	5.3
% Δ Layoff Rate	-100.0	-70.7	-4.6	44.7	-3.1
% Δ Avg. Log Real Wage	-2.3	1.4	-0.5	0.2	-0.5
Avg. Log Real Wage Growth (Stayers)	8.2	5.3	4.8	5.0	4.8
Avg. Log Real Wage Growth (Switchers)	18.9	11.9	10.4	10.5	10.4
Δ Unempl. Rate	-0.3	-0.2	-0.4	-0.1	-0.4

Notes: This table compares the effects of different shocks on the labor market. See text for additional details.

Our key finding in this section is that each of these other shocks is inconsistent with the broad set of observed labor market features during the 2021-2024 period. The results are summarized in Table 5. As with most "hot" labor market stories, all of these other shocks cause a large increase in the job-finding rate (the U-E rate). When the overall labor market is strong, unemployed workers have an easier time finding a job. However, as seen from Panel B of Figure 3, the U-E rate did not increase during this period. A key insight from this exercise is that most shocks that cause a large increase in the vacancy-to-unemployment rate also cause a sharp increase in the U-E rate. Moreover, it is well documented that the job-finding rate of the unemployed is very cyclically sensitive.[37] Second, these other shocks struggle to match the large decline in real wages during this period even though we are still imposing the same nominal wage rigidities. The positive productivity shock actually causes real wages to rise while the other three shocks result in little movement in real wages. Again, traditional hot labor market shocks that cause the vacancy-to-unemployment rate to increase struggle to explain the sharp decline in real wages observed during this period. Third, none of these shocks can match the large increase in the E-E rate seen in the data. Finally, all of these shocks fail to generate a large rise in the real wage growth of job changers relative to job-stayers.[38] Collectively, these other stories that can generate a rapid increase in the vacancy-to-unemployment rate fail to match other labor market flows and overall wage dynamics observed during the 2021-2024 period in the United States.

7.2. The Potential Causes of the Recent Inflation

There is growing evidence that the inflation observed in the United States between 2021 and 2023 was not caused by rising wages from an overheated labor market. For example, both Lorenzoni and Werning (2023b) and Bernanke and Blanchard (2024) provide evidence that the burst of inflation starting in mid-2021 in the United States was the result of shocks to prices holding wages fixed. One piece of evidence supporting their conclusion is that the large rise in aggregate prices predated the modest nominal wage increase. Instead, these authors conclude that the observed inflation resulted from some combination of (i) restricted aggregate supply coming from energy price increases, sectoral reallocation, and pandemic-induced supply constraints and (ii) increased aggregate demand coming from the large stimulus enacted during the pandemic. As we discuss next, these two shocks have opposite effects on firm labor demand.

Rising oil prices and supply chain backlogs due to pandemic closures have similar effects on the labor market as a negative aggregate productivity shock. These negative supply shocks will reduce labor demand given that firms will want to hire less labor. This will put

[37] See, for example, Elsby, Hobijn, and Sahin (2010) and Shimer (2012).

[38] Recent work by Bagga, Mann, Sahin, and Violante (2023) highlights how the shifting ability to work from home can generate many of the labor market patterns during the 2020s. The increasing E-E flows in their paper come from workers sorting to jobs that offer work-from-home opportunities. Given that working from home is an amenity to workers, their model predicts that the wages of job-changers should *fall* if they are seeking work-from-home opportunities. Nonetheless, the fact that real wages of switchers *grew* faster than stayers (Figure 5) implies that our mechanism is another key driver of labor dynamics during this period.

downward pressure on the vacancy-to-unemployment rate, EE flows, UE flows, vacancies, employment and average real wage growth (the opposite of the results in column 2 of Table 5). A negative supply shock would not generate a hot labor market. Conversely, a positive aggregate demand shock due to increased government spending or pent-up demand from the Pandemic would increase the demand for labor. This would have traditional hot labor market effects of rising the V-U ratio, U-E flows, vacancies, employment and real wages (similar to columns 3 and 4 of Table 5). These two shocks at the center of explanations for the current inflation have offsetting effects on labor demand. This could be a possible explanation for why aggregate employment (and GDP) did not change much during the current inflation period. If that is the case, the effects of inflation itself could be the primary driver of the real wage dynamics and labor market flows observed during the 2021-2024 period. As seen in the prior section, prior periods of aggregate supply shocks (the early-1950s, the mid-1970s and the late-1970s) had similar labor market dynamics.

8 Conclusion

The dramatic recent increase in the vacancy-to-unemployment rate has renewed interest among both academics and policymakers about the causal effect of tight labor markets on inflation. In this paper, we develop a model that combines elements of modern frictional labor markets with nominal wage rigidities to show that the causation can go in the opposite direction with high inflation driving a rise in the vacancy-to-unemployment rate giving the appearance of a tight labor market even as real wages fell. Calibrating the model with pre-2020 data, we show it matches well trends in worker flows and wage changes during the 2021-2024 period. We show evidence of our model mechanisms using historical data. In particular, prior periods of high inflation within the United States were associated with increases in vacancies and an upward shift in the Beveridge curve. We show that other shocks that can generate a rise in the vacancy-to-unemployment rate are unable to match key labor market statistics during the recent period. In particular, the other hot labor market shocks we explore struggle to generate falling real wages, a constant job-finding rate of the unemployed, the sharp rise in E-E flows, a sharp decline in layoffs, and the large increase in wages of job-changers relative to job-stayers.

Specifically, our model incorporates nominal wage rigidities into a framework with heterogeneous workers and frictional labor markets with many types of endogenous worker flows (quits, layoffs, and on-the-job search). The nominal wage rigidities along with a two-sided lack of commitment on the part of workers and firms results in inefficient separations initiated by both workers and the firms. In this environment, a burst of inflation, all else equal, reduces real wages on impact. In order to have their real wage keep up with inflation, workers can search for another job where they could contract over an updated real wage or they could try to take steps to renegotiate their wage with their existing firm. Costly worker search effort and sticky wages imply that a burst of inflation reduces workers real wages and incentivizes job-to-job transitions that lead to a rise in aggregate vacancies. Our quantitative model explains why the Beveridge Curve steepened during the recent inflation period. In particular, we illustrate how inflation can cause a large rise in vacancies—through increased job-to-job transitions—creating the appearance of a tight labor market without any additional labor market shocks.

We estimate that the recent burst of inflation reduced the welfare of all workers throughout the wage distribution; however, the losses were greatest for high-wage workers. Specifically, we find workers in the bottom, median, and top wage deciles experienced a decline in welfare of roughly 75%, 85%, and 110% of their monthly real income, respectively. The welfare costs come from the fact that real wages initially declined and from the fact that workers subsequently took costly actions to offset the initial real wage losses. Conversely, we also find that inflation had a positive effect on worker well-being by reducing layoffs. On net, the predominant reason for the welfare loss of workers is because real wages declined; the costs of additional search and renegotiation costs were roughly offset by the gains from the reduced layoff margin. Collectively, our findings provide a model-based rationale for why workers currently report that they dislike inflation in the current labor market environment.

Our welfare results, however, are likely incomplete. In our model, we assumed firms are homogeneous. In a world with heterogeneous firms and match specific productivity between workers and firms, there could be an additional productive role for job-to-job flows in that workers could have a steeper job ladder. In this case, the extra search induced by a burst of inflation could potentially increase aggregate productivity by incentivizing workers to more quickly move up their job ladder. Incorporating heterogeneous firms and match specific productivity into our framework would be a fruitful area of future work to get a more complete picture of the potential welfare effects of inflation through the labor market.

Finally, we show that the inflation redistributes match surplus from workers towards firms.

The decline in real wages for workers is matched by an increase in profits on the part of firms. In this case, our model can also explain why the ratio of corporate profits to GDP was at historically high levels during the 2021-2023 inflation period.

References

Afrouzi, H., A. Dietrich, K. Myrseth, R. Priftis, and R. Schoenle (2024): "Inflation preferences," Discussion paper, National Bureau of Economic Research.

Albertini, J., A. Poirier, and D. Trupkin (2019): "A Job Vacancy Rate for Argentina," Discussion paper, GATE Lyon Saint-Etienne Working Paper 1919.

Alves, F., and G. L. Violante (2023): "Some Like It Hot: Monetary Policy Under Okun's Hypothesis," .

Autor, D., A. Dube, and A. McGrew (2024): "The Unexpected Compression: Competition at Work in the Low Wage Labor Market," Working Paper 31010, National Bureau of Economic Research.

Bagga, S., L. Mann, A. Sahin, and G. Violante (2023): "Job Amenity Shocks and Labor Reallocation," Discussion paper, Working Paper.

Ball, L., D. Leigh, and P. Mishra (2022): "Understanding US Inflation during the COVID-19 Era," *Brookings Papers on Economic Activity*, 53(2 (Fall)), 1–80.

Barnichon, R. (2010): "Building a composite Help-Wanted Index," *Economics Letters*, 109(3), 175–178.

Barro, R. J. (1972): "A theory of monopolistic price adjustment," *The Review of Economic Studies*, 39(1), 17–26.

Benigno, P., and G. B. Eggertsson (2023): "It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve," Working Paper 31197, National Bureau of Economic Research.

Bernanke, B., and O. Blanchard (2024): "What Caused the U.S. Pandemic Era Inflation?," *American Economic Journal: Macroeconomics*.

Bils, M., M. Kudlyak, and P. Lins (2023): "The Quality-Adjusted Cyclical Price of Labor," *Journal of Labor Economics*, 41(S1), S13–S59.

Blanco, A., and A. Drenik (2023): "How Does Inflation "Grease the Wheels" in a Frictional Labor Market?," Manuscript.

Blanco, A., A. Drenik, C. Moser, and E. Zaratiegui (2024): "A Theory of Labor Markets with Inefficient Turnover," Working Paper 32409, National Bureau of Economic Research.

Chan, M., E. Mattana, S. Salgado, and M. Xu (2023): "Wage Setting and Passthrough: The Role of Market Power, Production Technology and Adjustment Costs," Discussion paper, Working Paper.

Cheremukhin, A., and P. Restrepo-Echavarria (2023): "The Dual Beveridge Curve," Manuscript.

Chodorow-Reich, G., and L. Karabarbounis (2016a): "The Cyclicality of the Opportunity Cost of Employment," *Journal of Political Economy*, 124(6), 1563–1618.

——— (2016b): "The cyclicality of the opportunity cost of employment," *Journal of Political Economy*, 124(6), 1563–1618.

Davis, S. J., R. J. Faberman, and J. C. Haltiwanger (2013): "The establishment-level behavior of vacancies and hiring," *The Quarterly Journal of Economics*, 128(2), 581–622.

Durante, A., J. Larrimore, C. Park, and A. Tranfaglia (2017): "Report on the economic well-being of us households in 2016," Discussion paper, Board of Governors of the Federal Reserve System (US).

Ellieroth, K., and A. Michaud (2024): "Quits, Layoffs, and Labor Supply," Opportunity and Inclusive Growth Institute Working Paper #94.

Elsby, M. W. L., B. Hobijn, and A. Sahin (2010): "The Labor Market in the Great Recession," *Brookings Papers on Economic Activity*, 41(1 (Spring), 1–69.

Fujita, S., G. Moscarini, and F. Postel-Vinay (2024): "Measuring Employer-to-Employer Reallocation," *American Economic Journal: Macroeconomics*, 16(3), 1–51.

Galí, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gertler, M., C. Huckfeldt, and A. Trigari (2020): "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," *The Review of Economic Studies*, 87(4), 1876–1914.

Grigsby, J., E. Hurst, and A. Yildirmaz (2021): "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, 111(2), 428–71.

Guerreiro, J., J. Hazell, C. Lian, and C. Patterson (2024): "Why Do Workers Dislike Inflation? Wage Erosion and Conflict Costs," Discussion paper, Working Paper.

Hajdini, I., E. S. Knotek, J. Leer, M. Pedemonte, R. W. Rich, and R. Schoenle (2022): "Low Passthrough from Inflation Expectations to Income Growth Expectations: Why People Dislike Inflation," Working papers, Federal Reserve Bank of Cleveland.

Hall, R. E. (2017): "High discounts and high unemployment," *American Economic Review*, 107(2), 305–330.

Hazell, J., and B. Taska (2024): "Downward Rigidity in the Wage for New Hires," *American Economic Review*, forthcoming.

Hosios, A. J. (1990): "On the efficiency of matching and related models of search and unemployment," *The Review of Economic Studies*, 57(2), 279–298.

Hyatt, H. R. (2015): "The decline in job-to-job flows," *IZA World of Labor*, pp. 175–175.

Jarosch, G. (2023): "Searching for job security and the consequences of job loss," *Econometrica*, 91(3), 903–942.

Kudlyak, M. (2014): "The cyclicality of the user cost of labor," *Journal of Monetary Economics*, 68, 53–67.

Lorenzoni, G., and I. Werning (2023a): "Inflation is Conflict," Manuscript.

——— (2023b): "Wage-Price Spirals," *Brookings Papers on Economic Activity*, Fall.

Moscarini, G., and F. Postel-Vinay (2023): "The Job Ladder: Inflation vs. Reallocation," Working Paper 31466, National Bureau of Economic Research.

Petrongolo, B., and C. A. Pissarides (2001): "Looking into the black box: A survey of the matching function," *Journal of Economic literature*, 39(2), 390–431.

Pilossoph, L., and J. Ryngaert (2023): "Job Search, Wages, and Inflation," Discussion paper, Working Paper.

Pilossoph, L., J. Ryngaert, and J. Wedewer (2024): "The Search Costs of Inflation," Discussion paper, Mimeo-Duke University.

Reed, S. B. (2014): "One hundred years of price change: the Consumer Price Index and the American Inflation Experience," *Monthly Labor Review*, U.S. Bureau of Labor Statistics.

Sheshinski, E., and Y. Weiss (1977): "Inflation and costs of price adjustment," *The Review of Economic Studies*, 44(2), 287–303.

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

——— (2012): "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 15(2), 127–148.

Stantcheva, S. (2024): "Why Do We Dislike Inflation?," *Brookings Papers on Economic Activity*, Spring.

Volpe, O. (2024): "Job Preferences, Labor Market Power, and Inequality," Discussion paper, Working Paper.

Online Appendix

A Data Description

In this section of the appendix, we discuss in detail the data work we use in the paper.

A.1. JOLTS

We use the Job Openings and Labor Turnover Survey (JOLTS) data to measure quits, layoffs, and vacancies during the December 2001 through the June 2024 period. We downloaded the data directly from the JOLTS data website when creating the descriptive work shown in Section 2.[39] The JOLTS dataset, collected by the U.S. Bureau of Labor Statistics (BLS) provides a snapshot of worker hiring and separation flows for a nationally representative sample of non-farm business and government employers during a given month. Below, we provide definitions of the JOLTS Layoff Rate, Quit Rate, and Vacancy Rate.

Layoff Rate: The *layoff rate* reflects all workers who were involuntarily terminated by a firm during a given month divided by total monthly employment. Involuntary terminations include workers laid-off with no intent to rehire; workers fired or discharged for cause; workers whose discharge resulted from mergers, downsizing, or firm closings; and seasonal workers discharged at the end of the season.

Quit Rate: The *quit rate* reflects workers who left voluntarily during the month divided by total employment at the end of the month. The quit rate captures workers who left the firm by either (i) flowing into unemployment before starting to look for another job (a voluntary "E-U" flow), (ii) directly transitioning to another firm (an "E-E" flow), or (iii) leaving the labor force (an "E-N" flow). Panel A of Appendix Figure A.1 highlights the close relationship between the monthly CPI inflation rate and the monthly quit rate during this period. Each observation in the figure is a month during the 2016 to 2024 period. As seen from Panel A, there is a strong positive relationship between monthly year-over-year price inflation and the monthly quit rate. A simple linear regression through the scatter plot finds that a 1 percentage point increase in the inflation rate is associated with a 0.103 percentage point increase in the quit rate (standard error = 0.007); the R-squared of the regression was 0.70.

Vacancy Rate: The *vacancy rate* (or job-opening rate) is the number of open positions on the last business day of the month divided by the sum of employment and vacancies on the last day of the month. This data was also used when making the vacancy-to-unemployment rate series shown in Panel A of Figure 1. Panel B of Appendix Figure A.1 shows the tight relationship between the year-over-year CPI inflation rate and the monthly vacancy rate over

[39] See, https://www.bls.gov/jlt/data.htm.

Figure A.1: Monthly Inflation vs Monthly Labor Market Flows

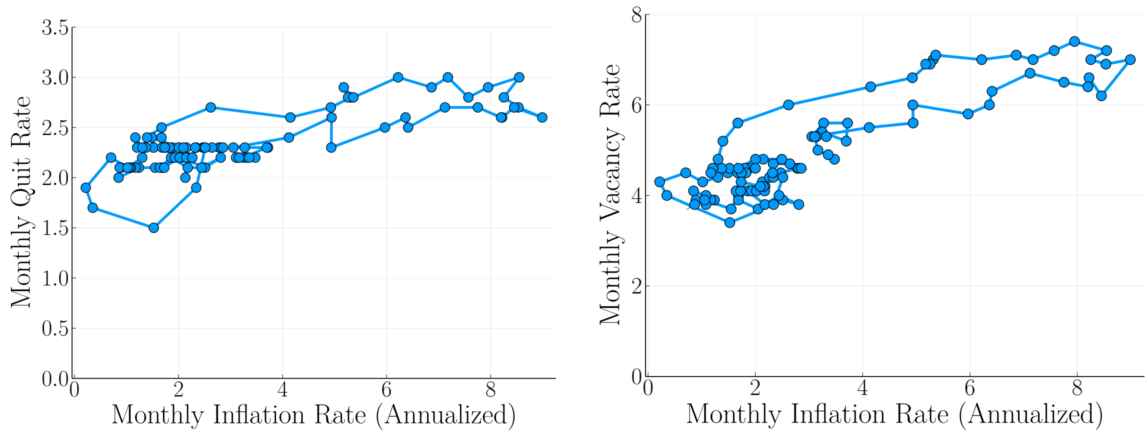

PANEL A: QUIT RATE VS. INFLATION RATE

PANEL B: VACANCY RATE VS. INFLATION RATE

Notes: Figure shows a scatter plot of the year-over-year CPI inflation rate vs the monthly quit rate (Panel A) and the monthly vacancy rate (Panel B). Each observation is a month between January 2016 and May 2024. The quit and vacancy rates are obtained from JOLTS while the inflation numbers are from the BLS'S CPI for urban consumers.

the entirety of the 2016 to 2024 period. This figure is analogous to the Beveridge curve but with the price inflation rate on the x-axis instead of the unemployment rate. While there has been a well-documented breakdown of the Beveridge curve during the last few years, the relationship between the inflation rate and the vacancy rate remained relatively stable during this time period. In particular, a simple linear regression through the scatter plot finds that a 1 percentage point increase in the inflation rate is associated with a 0.438 percentage point increase in the vacancy rate (standard error = 0.020); the R-squared of the regression was 0.83.

A.2. Atlanta Fed Wage Tracker

For our descriptive work on real wage growth during the 2016-2024 period, we use data from the Atlanta Fed Wage Tracker Index.[40] The Wage Tracker Index uses the panel component of the Current Population Survey (CPS) to make a measure of composition adjusted nominal wage growth. The structure of the CPS is such that individuals are in the sample for four months where they are surveyed about their labor market activities. After that, they leave the sample for eight months and then re-enter for a final four additional months. In their fourth

[40] We downloaded the data directly from https://www.atlantafed.org/chcs/wage-growth-tracker.

survey month and their eight survey month - which takes place one year apart - individuals are asked about their wages. The Atlanta Fed Wage tracker measures a year-over-year change in the workers per-hour nominal wage on their main job. For salaried workers, the hourly wage is computed as weekly earnings divided by usual weekly hours worked.

A.3. CPS Monthly Files

We use the Outgoing Rotation Group (ORG) of the Current Population Survey (CPS) to make our own measures of worker flows and wage dynamics across workers with different earnings. These moments are used to calibrate the worker heterogeneity within our model. In particular, we leverage this short panel to observe individual labor market flows by skill level (education) and initial earnings decile. Sample selection and estimation of labor market flows and earnings are described below.

A.3.1. Sample Selection. For the years 2015 to 2024, we select all individuals from ages 25 to 55, inclusive. We exclude all government employees, self-employed workers, and unpaid family members. This leaves us with 4,077,904 worker-month observations. This is the main sample we use in all the subsequent analyses.

A.3.2. Measure of Earnings. As noted above, in months 4 and 8 of the CPS interview waves, all employed workers in the week of the survey are asked about their usual weekly earnings. We use full-time workers' reported nominal weekly earnings as our wage measure. The CPS top codes individuals with nominal weekly earnings of higher than 2884 U.S. dollars. This top coding procedure changed in April 2023. To maintain consistency throughout our sample, we top code all workers with earnings of more than 2880 nominal U.S. dollars. Each month, we use the associated price index (CPI-U) published by the BLS to convert nominal earnings to real and create a cross-sectional real earnings distribution over time.

A.3.3. Labor Market Flows. The CPS basic monthly files report the employment status - employed or unemployed - of each individual in the labor force. We directly observe changes in employment status for each individual across adjacent months which provides us with a measure of gross flows from employment to unemployment (EU) and unemployment to employment (UE). In addition, we also observe a change in employers across adjacent months which allows us to measure EE flows. There are several technical issues that warrant discussion here. Fujita, Moscarini, and Postel-Vinay (2024) show that the CPS systematically underestimates EE flows since 2007 due to changes in survey methodology which induce selection on both unobservable and observable worker characteristics that are correlated with EE transitions. We use their published aggregate EE series to discipline our EE rates by earnings deciles and education group. Our raw EE estimates by various groups underestimate true EE flows, so we use a constant scaling factor to scale our EE flows by deciles to hit the aggregate EE rate as calculated by Fujita, Moscarini, and Postel-Vinay (2024). We use the following equation to determine our scaling factor:

$$EE = \frac{\alpha}{10} \sum_{d=1}^{10} EE_d \tag{A.1}$$

The decision of a constant scaling factor across deciles warrants discussion. If the elasticity of E-E probability varies with earnings then the scaling factor should be different. According to both Autor, Dube, and McGrew (2024) and our model, returns to search effort, in expectations, are decreasing in current earnings. Therefore, search effort, and in turn, EE rates are more elastic at the bottom of the earnings distribution which implies that a constant scaling factor underestimates the true EE rate for low earners and overestimates EE rates for high earners. Thus, our estimates which show that EE rates increased more for low earners relative to high earners is a conservative lower bound. Similar issues persist with estimating EU and UE flows using microdata in the CPS. We follow the seminal work of Shimer (2005) to infer the EU and UE using aggregate gross worker flows. Of course, data is observed in discrete time at a monthly frequency so we estimate the job finding rate (UE rate) by:

$$\lambda_t = 1 - \frac{U_{t+1} - U^s_{t+1}}{U_t} \tag{A.2}$$

where U represents the stock of unemployed workers at a given point in time and U^s represents the stock of short-term unemployed workers (unemployed for $<=$ 4 weeks). The separation rate (EU) rate is estimated using:

$$\delta_t = \frac{U^s_{t+1}}{E_t(1 - \frac{1}{2}\lambda_t)} \tag{A.3}$$

where E is the stock of employed workers at a given point in time.

B Additional Descriptive Results

In this section of the appendix, we show additional results as referenced throughout the main paper.

B.1. Evolution of Real Wages

We begin by showing that the real wage dynamics during the inflation period shown in Figures 1 and 4 are robust to alternative assumed real wage trends. In the main text, we constructed counterfactual real wages assuming they evolved according to the real wage trends during the pre-inflation ($2016 - 2019$) period. Figure B.1 shows that roughly similar gaps between expected and realized real wages emerge if we use the longer $2000 - 2019$ period to define our pre-period trend. These patterns are shown across the five panels in Figure B.1. In June 2024, nearly all groups remain below trend - with the median worker still 3.4% below their expected wage. The bottom income quartile, however, is now slightly above their expected wage in June 2024 when using the longer period to calculate the predicted trend; the trend in real wages for the bottom quartile worker was lower during the $2000-2019$ period than it was during the $2016-2019$ period.

B.2. Dynamics of Unemployment

Figure B.2 shows that the decline of unemployment starting in 2021 was largely driven by declining job destruction rate (or layoffs) rather than increasing job finding rate. We show that the decline in job destruction rates predicted by our theory are largely responsible for the unemployment dynamics since 2021. To see this easily, observe that if fluctuation in job-finding (red) explained all of the variation in unemployment (black) then the the two lines would be perfectly on top of each other, but clearly this is not the case. The fluctuations in job-finding predict persistently higher unemployment than what is observed. Changes in job destruction rate closely track the dynamics of unemployment during this period. This is a unique feature of labor market flows during the $2021-2024$ inflation period as Shimer (2012) finds that job-finding (rather than job-destruction) explains 80% of the variation in unemployment since 1948 in the US data. These results are consistent with the fact that the U-E rate did not change much during the inflation period; instead it was the decline in layoffs that were driving unemployment dynamics.

B.3. Duration of Vacancy

B.3 shows that the average time to fill a vacancy rose from about 30 days in the pre-period to 45 days during the peak of inflation. We use data on hires and vacancies at a monthly frequency from JOLTS to estimate the job-filling rate and back out the expected duration to fill a vacancy. Following the methodology described in Davis, Faberman, and Haltiwanger (2013), we assume that hires on day s of month t is given by:

$$h_{s,t} = f_t v_{s-1,t} \tag{B.1}$$

f_t is the daily job-filling rate which is constant over a given month, and $v_{s-1,t}$ is the stock of vacancies on day $s-1$ of month t. The above equations implies that a constant fraction f_t of vacancies are filled by new hires each day. Since data is reported at a monthly frequency, let $H_t = \sum_{s=1}^{26} h_{s,t}$. Then, in the steady state, the daily job-filling rate is given by:

$$f = \frac{H}{v}\left(\frac{1}{\tau}\right) \tag{B.2}$$

$\tau = 26$ represents the number of working days in a month. Given monthly data on hires and vacancies, the job-filling rate f_t can be directly estimated. The duration of a vacancy, in expectation, is given by $\frac{1}{f_t}$, the object of B.3.

Figure B.1: Evolution of Real Wages

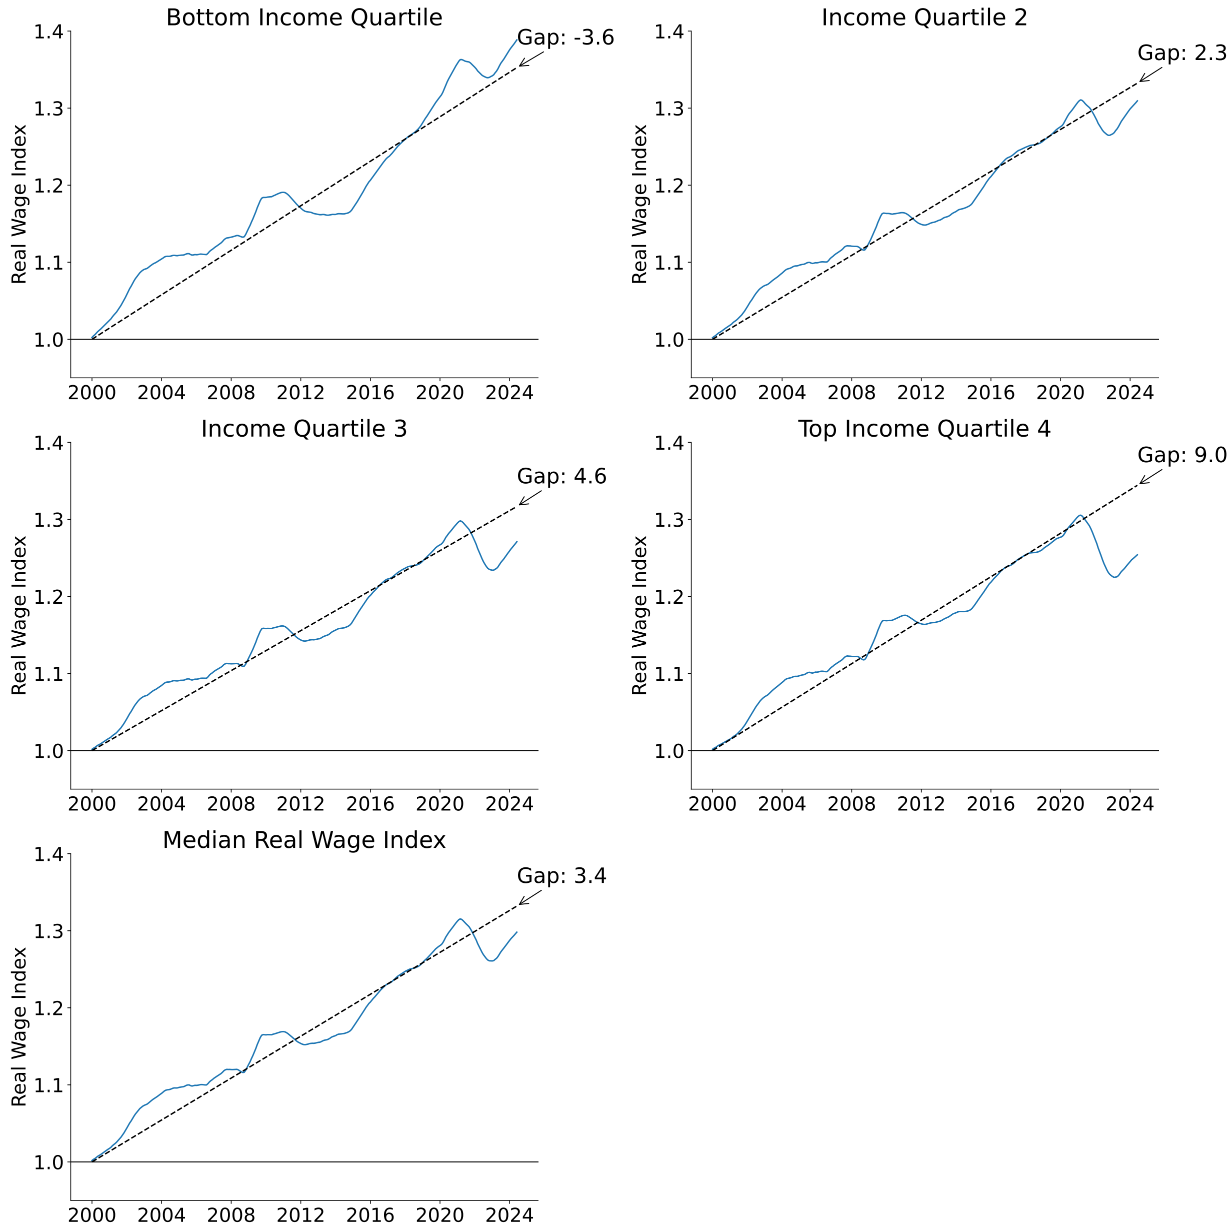

Notes: This figure shows the evolution of real wages across the income distribution between 2000 and 2024 for the same gropus shown in panel B of Figure 1 and in all panels of Figure 4. The blue line indicates realized real wages and the dotted black line shows the trend in real wage. The trend in wages is recovered separately for each group on wage data between 2000 – 2019. We project this trend on 2020 – 2024 data to create a counterfactual wage gap between predicted and observed real wage for each income group in June of 2024. The data is directly taken from Atlanta Fed Wage tracker.

Figure B.2: Decomposition of Unemployment Dynamics

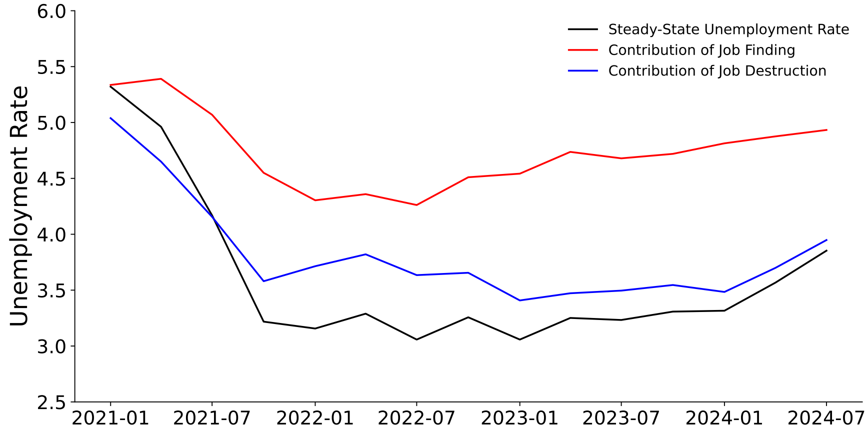

Notes: Contribution of fluctuations in the job finding (UE) and job destruction (EU) rates to fluctuation in the unemployment rate, 2021-2024, quarterly average of monthly data. The red line shows the counterfactual unemployment rate if all fluctuation were due to changes in the job finding rate ($\frac{\bar{\delta}}{\bar{\delta}+\lambda_t}$) and the blue line shows the counterfactual unemployment rate with only fluctuations in the job destruction rate ($\frac{\delta_t}{\delta_t+\bar{\lambda}}$). $\bar{\delta}$ and $\bar{\lambda}$ is the average job destruction and job finding rate between 1990 and 2024. The black line is the implied steady state unemployment rate ($\frac{\delta_t}{\delta_t+\lambda_t}$). This is a good approximation to the observed unemployment rate - correlation of .95 over $1990-2024$.

Figure B.3: Duration of Vacancy

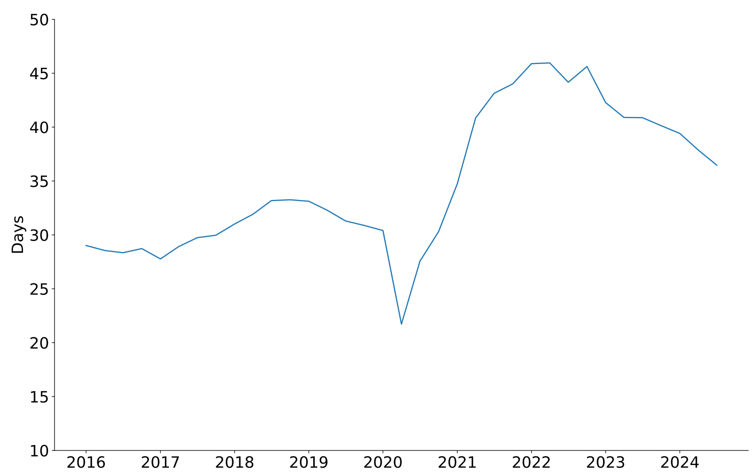

Notes: We estimate the job-filling rate given the data on the flow of hires and the stock of vacancies (see Davis, Faberman, and Haltiwanger (2013) for details). We take a quarterly average of the monthly job-filling rate and plot plot the implied duration of a vacancy.

B.4. E-U Flows, CPS Data

In this section, we show the E-U rate from the CPS. In particular, we downloaded the series “labor force flows employed to unemployed” and “all employees, total nonfarm” directly from the St. Louis Federal Reserves Economic Database (FRED) who extracted the series from CPS aggregates published by the BLS to make the E-U rate. For readability, we exclude the data from April 2020 from the graph when the EU rate exceeded 13%. As seen from the figure, the E-U rate did not change at all during the inflation period relative to the pre-period. All of the documented quits from the JOLTS data are showing up as increased E-E churn.

Figure B.4: E-U, CPS Data

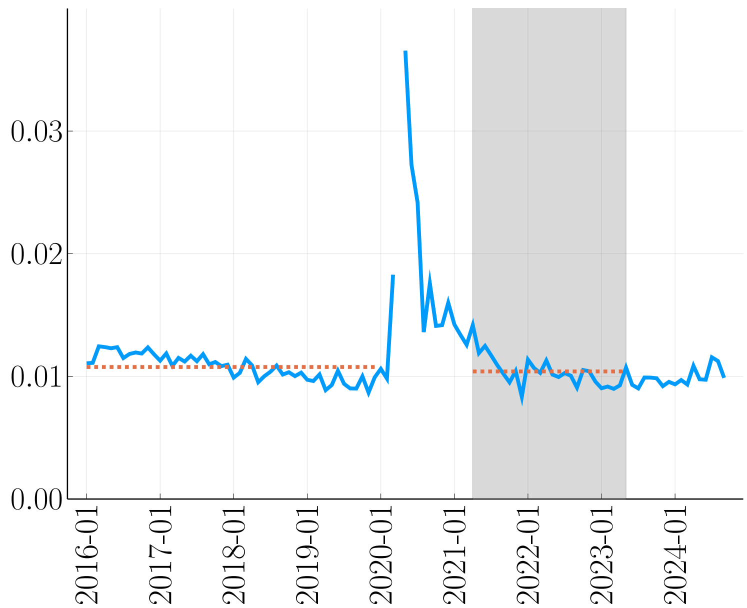

Notes: Figure shows the monthly E-U rate during the 2016-2024 period. Data downloaded directly from the St. Louis Federal Reserves Economic Database (FRED). In particular, we downloaded their series "labor force flows employed to unemployed" and "all employees, total nonfarm" to make the E-U rate. For readability, we exclude the data from April 2020 from the graph when the EU rate exceeded 13%.

B.5. E-E Flows, By Education Group

Appendix Table B.1 shows the change in the E-E rate during the pre-inflation period and the September 2021 through December 2022 period by education group.[41] We use education group as a proxy for the individual's position in the income distribution. These data are the same as we used for the aggregate E-E rate series in the main paper. The figure documents that the change in the E-E rate was not constant throughout the income distribution. Lower educated (lower income) individuals had a much larger change in the E-E rate (0.23 percentage points) than did higher educated (higher income) individuals (0.11 percentage points).

[41]We highlight the change in the E-E rate during a subset of the inflation period given that, as seen in Figure 3, this is the period when E-E rates were the highest.

Table B.1: Change in E-E Flows by Education during Inflation Period

Education	2016M1-2019M12	2021M9-2022M12	Change
Less than Bachelors	2.34%	2.57%	0.23 p.p. (0.04)
Bachelors or More	2.22%	2.33%	0.11 p.p. (0.05)

Notes: The table shows the average E-E rate for individuals with less than a Bachelor's degree (top row) and individuals with a Bachelor's degree or more (bottom row) during the inflation pre-period (column 1) and then again during the September 2021 through December 2022 period (column 2). Column three shows the difference in E-E rates between the two periods with the standard error of the difference in parentheses. The sample is restricted to those aged 25-55 from the monthly CPS files.

B.6. Job-Stayers vs Job-Changer Wage Growth, Atlanta Fed

The Atlanta Fed Wage Tracker Index also measures the nominal wage growth of job-stayers relative to job-changers over time. The underlying data for the Atlanta Fed Wage Tracker comes from the CPS. A key limitation of using the CPS data to measure the wage growth of job-changers is that the CPS follows addresses not people. If someone moves addresses they drop out of the CPS. Job-changers - particularly those that get large wage increases - are more likely to move locations than job-stayers. So, the CPS data may be downward biased for the wage growth of job-changers because the data does not capture the large wage changes of job-changers who move. None-the-less, the patterns in the Atlanta Fed data are broadly similar to what we show in the main paper using ADP data. In particular, as seen in Appendix Figure B.5, the gap in wage growth between job-changers and job-stayers doubled during the inflation period–just like the doubling observed in the ADP data. However, relative to the ADP data, the wage growth of job-changers relative to job-stayers is smaller in levels both during the pre-period and the inflation period consistent with the fact that the CPS may be missing some of the big wage changes of job-changers who also change residences.[42]

B.7. Corporate Profits

Appendix Figure B.6 shows the corporate profit to GDP ratio in the United States between 2016 and 2024 (quarterly). We downloaded this data directly from the FRED website. In particular, we used the series Corporate Profits After Tax (without IVA and CCA Adjustments)

[42]There are other differences between the ADP data and the CPS data in terms of the wage change measures of job-stayers vs job-changers. For example, the ADP Wage Tracker Index also includes signing bonuses and other forms of income in their wage series for job-stayers and job-changers.

Figure B.5: Nominal Wage Changes of Job-Stayers vs Job-Changers: Atlanta Fed Wage Tracker

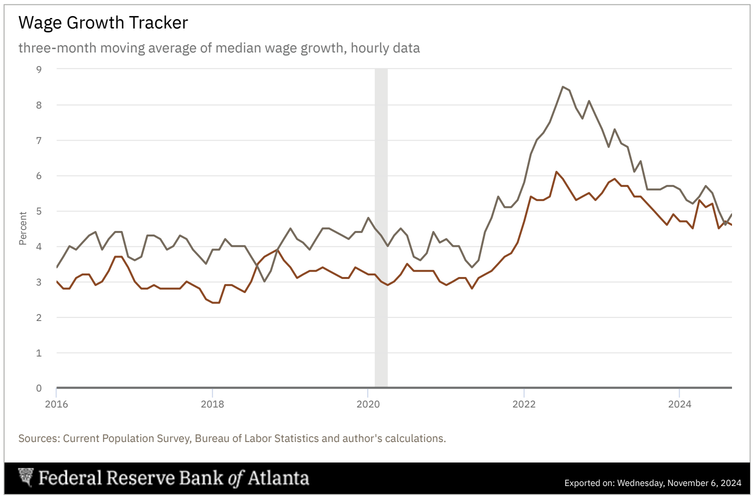

Notes: Figure shows nominal wage growth of job-stayers (red bottom line) and job-changers (grey top line) from the Atlanta Fed's Wage Tracker Index. See text for additional details of the data. We downloaded this figure directly from the Atlanat Fed's Wage Tracker website.

and divided that series by US Nominal GDP. As seen from the figure, the corporate profit to GDP ratio jumped from about 10% in the 2016-2019 period to 11.6% during the inflation period. The corporate profit to GDP ratio during the inflation period is the highest it has been since 1950. Between 1950 and 2020, there were only 9 quarters where the corporate profit to GDP ratio exceeded 11% and there were no quarters where the ratio exceeded 12%. The current corporate profit to GDP ratio is at historically high levels. It should be noted that the corporate profit to GDP ratio was also at historically high levels in 1950, 1974, and 1979 – all periods where the both the inflation rate was high and the labor market was not particularly strong. Specifically, between 1950 and 2000, there were only four periods when the corporate profit to GDP ratio exceeded 7%; three of those were the early 1950s, 1974, and 1979.

The rise in the corporate profit to GDP ratio is consistent with the prediction of our model where firm labor market power increased during the inflationary period because nominal wages are sticky. The rise in the corporate profit to GDP ratio at face value is inconsistent with other theories suggesting firm labor market power decreased during the post-pandemic period due to the labor market being tight.

Figure B.6: Corporate Profits to GDP Ratio

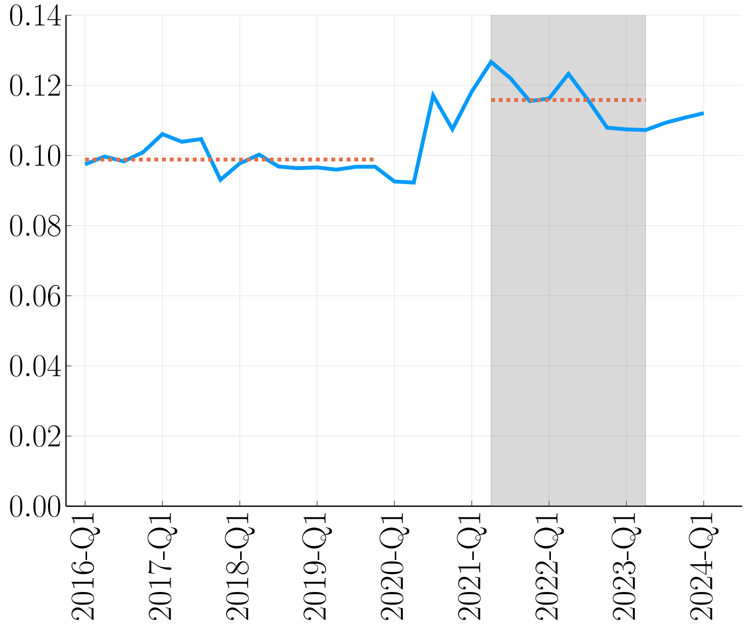

Notes: Figure shows the U.S. corporate profits (after tax, without inventory valuation adjustment and capital consumption adjustment) relative to nominal GDP. Data from the U.S. Bureau of Economic Analysis retrieved from FRED, Federal Reserve Bank of St. Louis.

C Alternate Mechanism Analysis

In this section, we discuss the procedure for exploring other shocks that can generate a rising vacancy-to-unemployment rate through the lens of our model. In particular, we define the size of the various other shocks we explore so they roughly match our baseline increase in the vacancy-to-unemployment rate on impact from the one-time 13.5% increase in inflation (shown in Panel A of Figure 15); one impact of our baseline one-time inflation shock, the vacancy-to-unemployment rate increased by 8.5%. Specifically, we explore four different shocks: a positive shock to aggregate productivity (A_t), a negative shock to the household discount rate (ρ), a negative shock to the level of the vacancy posting cost (K), and a negative shock to the value of non-employment (B). We calibrate the model so all four of these shocks increase vacancies relative to unemployment at roughly the same magnitudes observed in our baseline model of the one-time inflation increase. Throughout these additional exercises, we still impose nominal wage rigidities.

Appendix Table C.1 shows the results of these exercises. Across the columns are the different shocks. In the first column is our baseline shock of a one-time monetary expansion such that it increases the price level by 13.5% on impact. These baseline results are described in the first set of counterfactuals in Section 5. In columns 2-5, respectively, we show the results for the one-time unexpected productivity increase (A_t), lower discount rate (ρ), lower vacancy posting cost (K) and lower non-employment benefit (B). The rows highlight various labor market outcomes. The first four rows measure how various labor market flows respond upon: the vacancy to unemployment rate (row 1), the percentage change in E-E flows (row 2), the percentage change in U-E flows (row 3) and the percentage change in the layoff rate (row 4). All five shocks – by design – match the increase in the vacancy-to-unemployment rate; the shock sizes were chosen to match the roughly 8.5 percentage point increase in the vacancy to unemployment rate. Key features of the data during the recent inflation period were that the E-E rate increased sharply, the U-E rate hardly changed, and the layoff rate fell. In order to match the rising vacancy-to-unemployment rate, all of the other generate a relatively large increase in the U-E rate and very small increases in the E-E rate. Normal hot labor market shocks that causes measured market tightness to increase do so by making the returns to working increase relative to the returns to staying at home implying a larger increase in the U-E rate. Notice, only the large productivity shock results in a large decline in layoffs. In our baseline model, the large inflation reduces workers' real wages relative to their productivity; this makes workers relatively cheap from the firm's perspective reducing their desire to layoff workers. With the productivity shock and sticky wages, the story is similar. The positive productivity shock reduces worker real wages relative to their productivity making them cheap from the firm's perspective.

Rows 5-7 show the log change in real wages after 12 months, the change in per-period real wage growth on impact for job stayers, and the change in per-period real wage growth for job changers. In the data, during the inflation period real wages fell and the gap in the growth rate of job-changers grew sharply relative to job-stayers. Only our baseline shock can generate a large decline in real wages while simultaneously generating an increase in the vacancy-to-unemployment rate. The productivity shock increases real wages one year out. The other three shocks have little effect on real wages. Importantly, only the inflation shock in column 1 can generate a sharp increase in the real wage growth of job-stayers relative to job-changers. In steady state, the real wage growth of job-changers is about 5 percentage points higher than job-stayers. For the other shocks, this 5 percentage point gap remained. It is only with the baseline shock that the gap in wage growth between job-stayers and job-changers widened.[43] As the table shows, the only shock that matches the rising in the vacancy-to-unemployment rate will simultaneously matching other labor market outcomes is our baseline shock.

[43]The last row shows percentage point change in the unemployment rate after one-year. This change is similar across all the shocks.

Table C.1: Comparison of alternative mechanisms

Variable	Baseline	Higher Agg. TFP	Lower ρ	Lower K	Lower B
Δ V/U Ratio	8.4	8.4	8.7	8.7	8.4
% Δ EE Rate	41.3	7.5	2.9	5.3	2.6
% Δ UE Rate	0.3	2.9	5.5	4.2	5.3
% Δ Layoff Rate	-100.0	-70.7	-4.6	44.7	-3.1
% Δ Avg. Log Real Wage	-2.3	1.4	-0.5	0.2	-0.5
Avg. Log Real Wage Growth (Stayers)	8.2	5.3	4.8	5.0	4.8
Avg. Log Real Wage Growth (Switchers)	18.9	11.9	10.4	10.5	10.4
Δ Unempl. Rate	-0.3	-0.2	-0.4	-0.1	-0.4

Notes: This table compares the effects of different shocks on the labor market. See text for additional details.