Job Search Under Low Pay, Low Security, and High Unemployment

Scott Abrahams

October 21, 2021

[Click here for the most recent version]

Abstract

The post-Covid shutdown U.S. labor market is characterized by many open jobs and substantial unemployment. Despite numerous hypotheses regarding this paradoxical situation, existing models in labor economics lack the features necessary to adequately evaluate the potential policy remedies or causes. I construct a search model to reproduce the prevailing dynamics, which I then use to evaluate the employment impact and effect on the welfare of low-income workers from large-scale transfers, unemployment benefits, and an earnings subsidy. One-off transfers hardly affect employment in a downturn. Raising unemployment benefits by $300 per month leads to a 0.6 point increase in the unemployment rate. Adding an earnings subsidy to unemployment benefits raises employment by 7 percentage points at a 23% lower level of inequality, yet increases costs by only 16-30%. In contrast to a compensating differentials framework, the model shows that lower earners tend to be more at risk for layoffs, and that this effect is stronger following a negative aggregate shock. Following a simulated shock, mean consumption falls by 5% yet consumption at the 10th percentile drops by 16%. Although the risk borne by lower earners is often greater, their savings are often lower, because they have less consumption to seek to smooth and less opportunity to accumulate savings.

1 Introduction

We are in the midst of an unusual labor market situation. Although the U.S. economy has reopened following the Covid-19 shutdown, it has not returned as expected. Instead, we have many unfilled job openings, many unemployed workers, and yet are somehow facing a significant labor shortage. As a leading labor economist put it in an opinion article, "We don't know what's going on" (Autor, 2021). Several theories abound that seek to explain the ongoing labor market phenomenon. Twenty-eight percent of economists surveyed believe that overly generous unemployment benefits are deterring workers from accepting low-wage offers, with another half reporting being uncertain.[1] Another theory posits that people have saved the emergency Covid relief payments and are now self-financing greater pickiness.[2] A third possibility offered by Autor (2021) is that people have begun to value their time not spent working more relative to the consumption potential from employment, in particular employment at very low paying jobs. Despite the numerous hypotheses, existing models in labor economics lack the features necessary to adequately evaluate the situation. In particular, these hypotheses imply an important distinction between income and consumption in regards to worker welfare, which is assumed away by the majority of models, as they do not allow for saving. They

[1] https://www.cnn.com/2021/06/03/politics/unemployment-benefits-worker-shortage/index.html
[2] https://www.theatlantic.com/ideas/archive/2021/06/labor-shortage-positive/619050/

also imply that the population of unemployed workers consists primarily of those who were formerly low earners. Investigating this situation requires a dynamic model of the labor market where workers can save, where exposure to job loss differs by income, and where an unexpected aggregate downturn can alter these calculations.

The purpose of this paper is improve our understanding of the environment and incentives facing workers in this type of scenario. With that objective, I construct a model of the labor market that more accurately represents these dynamics and which allows me to conduct realistic counterfactual experiments to evaluate potential policy remedies or causes. I pay particular attention to the welfare of low-income workers, as they are the most susceptible to the costs of a downturn, as I will show, and are often the people we have in mind when we think about inequality.

The main policy takeaway from the experiments is that combining unemployment benefits with an earnings subsidy for low-income workers may be a cost-effective way to support consumption in a downturn while incentivizing employment. As hypothesized, compared to no unemployment insurance, a $1,000 per month unemployment benefit does deter job acceptance, raising the unemployment rate by 3.3 points in normal times and up to 6.2 points during a downturn. Yet conditional on there being unemployment benefits, the difference in employment outcomes between policies of $1,000 and $1,300 is only 0.6 points after a year. In an alternate experiment, I augment the $1,000 unemployment benefit with a tapered earnings subsidy for low-income workers. Compared to unemployment insurance alone the combined policy raises the employment rate by 7 percentage points four quarters into a downturn at a 23% lower level of inequality as measured by the ratio of the 90th to 10th percentiles of consumption. Although the combined policy is by definition more expansive, because of the strong employment incentive total program costs are only 30% higher after one year and 16% after two, with the gap narrowing as the economy adjusts to the policy incentives. Costs are only 7% higher than the policy that raises unemployment benefits by $300.

The main takeaway from the model is that the dynamics of layoff risk and earnings can cause low earnings to be a persistent state because of increased churning into unemployment. Through the channel of unemployment, low earnings in one job raise the likelihood of also earning little in the job after that. The tendency is towards inequality, as higher earners also enjoy greater security as they move up the job ladder. Furthermore, this distributional job loss dynamic is amplified during an economic downturn, where mean consumption falls by 5% yet consumption at the 10th percentile drops by 16%. Although the risk borne by lower earners tends to be greater, their savings tend to be lower, because there is less consumption to want to smooth with savings and less opportunity to accumulate those savings. Counterfactual simulation shows that savings are still an important form of self-insurance, without which the predicted drop in mean consumption induced by a downturn is double in magnitude. However, raising everyone's savings by distributing $5,000 checks has little measurable impact on employment outcomes despite costing 600 times what a $1,000 per month unemployment benefit costs each month.

What type of worker is most likely to be laid off? Under a compensating differentials framework, jobs that entail greater layoff risk should be those that offer higher wages. However, in a frictional environment such differentials may not arise (Hwang et al. 1998, Bonhomme & Jolivet 2009). Empirical evidence, depicted in Figure 1, suggests that contrary to the compensating differentials framework, there is a negative relationship between layoff risk and earnings across the left tail of the earnings distribution followed by constant risk in the right tail. Using 2019 data from the Current Population Survey (CPS), the figure plots local polynomial regression results for the monthly transition rate into unemployment against weekly earnings for wage and

salary workers.[3] Job loss is most likely for the lowest earners. Layoff risk decreases as earnings rise, levelling off around $1000 per week, or just below median full-time earnings in the United States as of 2021.[4] Omitting this dynamic from a model would lead us to conclude incorrectly that the population of unemployed workers was drawn at random from across all previous income-levels.

Figure 1: Monthly transition rate into unemployment by weekly earnings

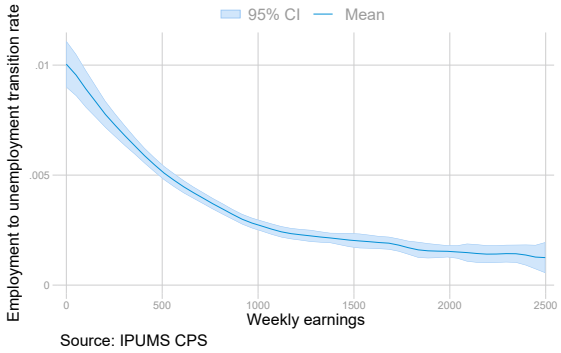

Moreover, this earnings-risk gradient grows even steeper during a downturn, as shown in Figure 2. The left-hand plot depicts the monthly transition rate into unemployment against weekly earnings in the boom of 2006 against those in 2008 amid the Great Recession. Layoff risk shifts upward by almost 50% for the lowest earners, from about 0.009 to 0.013, rises less for median earners, and is only slightly higher or unchanged for workers earning $1500 or more. The same gap appears in the plot on the right-hand side, which compares CPS data for 2019 and 2020. During the Covid-19 pandemic of 2020, layoff risk rises by about one percentage point for the lowest earners over the previous year, by half a percentage point for median earners, and only a quarter percentage point for higher earners. Despite certain tangentially related observations such as how cyclical variation in job displacement is higher for less educated workers (Farber, 2015), and how less productive firms are more likely to lay off workers in economic contractions (Haltiwanger et al., 2021), I have not seen a clear characterization to date of the pattern shown here regarding the relative increases in job loss rates during a downturn across earnings levels. A model that cannot reproduce this pattern would by construction assume that the impact of a downturn was even, when in fact we have strong evidence that lower earners are more likely to be affected.

[3] See Appendix F for data details and additional figures by subgroup to show that this pattern is not driven by composition effects.

[4] `https://www.bls.gov/news.release/pdf/wkyeng.pdf`

Figure 2: Layoff risk and earnings before and during two recent downturns

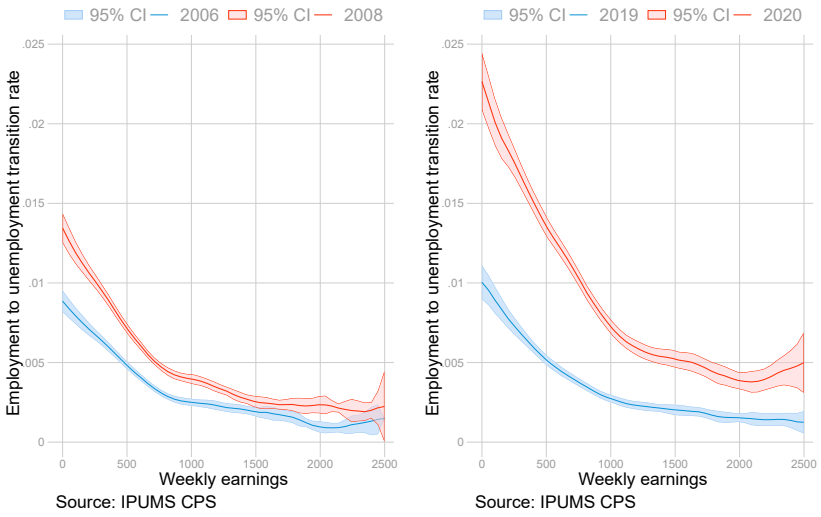

Taken together, these facts tell us something important about the environment facing low-income workers. First, individual preferences over the risk of job loss do not determine who is actually at risk. That is, workers do not choose an optimal bundle of risk and earnings. Rather, low risk and high earnings tend to go hand in hand, such that regardless of preferences there is a sense in which certain jobs are unambiguously more desirable than others. Second, a bad aggregate economic state has clear distributional consequences. While high-income workers may likely retain their positions, low-income workers lose theirs with greater frequency. Supposing that wage offers also fall during the negative state, the replacement jobs available may pay less than even the low-paying jobs from which those workers were displaced. In addition to informing our approach to understanding the Covid shock, these implications are consistent with what we observe following local labor market shocks, including the extensive literature on the earnings losses of displaced workers (Jacobson et al., 1993), Couch & Placzek (2010), and the higher unemployment rates, more frequent job churning, and lower lifetime earnings for workers in the areas most affected by the China Shock of the early 2000s (Autor et al., 2016).

The remainder of the paper proceeds as follows: Section 2 summarizes the literature relevant to the current work. Section 3 develops the model. Section 4 describes the data and sample construction. Section 5 details the estimation strategy. Section 5 reviews the results and evaluates model fit. Section 6 conducts the counterfactual analysis of a downturn and the policy experiments. Section 8 reviews the main conclusions and considers their implications. Additional details are provided in the appendices.

2 Related literature

The ideas in this paper are built on an extensive groundwork of previous economic research. In this section I characterize the related literature and explain how those studies inform my work. As the general topic of labor market supply and social insurance programs has been discussed extensively in other sources (see Krueger & Meyer (2002) and Moffitt (2002) for summaries), aside from a discussion at the end about a few specific sources, I focus here on the literature as it relates to my modeling choices.

Underpinning my modeling approach is the idea that there are widespread productivity differences across firms. Evidence for such differences is summarized in Syverson (2011). Two empirical estimates are Syverson (2004), who finds that the average difference in log total factor productivity between the 90th and 10th percentile of firms in the United States is 0.651, and for Hsieh & Klenow (2009), who estimate a ratio in levels of 5 to 1 in India and China. The U.S. Census Bureau and Bureau of Labor Statistics now measure the dispersion of intra-industry productivity with the Dispersion Statistics on Productivity (DiSP) series. Their initial statistics show that an establishment at the 75th percentile of within-industry productivity is more than twice as productive as one at the 25th percentile (Cunningham et al., 2019).

Given productivity differences among firms, my model claims that those differences relate to differences in wages and in survival. The relationship between wages and productivity receives particular attention in the literature seeking to explain wage dispersion. Card et al. (2018) summarize the evidence that changes in the dispersion of productivity correlate with changes in the dispersion of wages, characterizing estimates of the wage-productivity elasticity as falling in the range 0.05-0.15 and the importance of firm wage effects as explaining about 20% of the variance of wages. Lamadon et al. (2019) find that matches differ in the rents they produce, interpreted as reflecting some sense of productivity, which workers and firms split fairly evenly. Evidence of pass-through from productivity to wages is also reviewed in Taber & Vejlin (2016). Much of this work, following the two-way fixed effects approach of Abowd et al. (1999), is concerned with controlling for positive sorting between skilled workers and productive firms. I do not model skill heterogeneity because my primary focus in on firm behavior related to job destruction; however I do show that the main implications of the model are robust to that refinement. To do so I rely on the claim in Lamadon et al. (2019) that sorting is due to production complementarities, and the finding in Card et al. (2018) that firms offer the same premiums to workers of all skill levels.

What the literature on productivity and wage dispersion leaves somewhat open is the link between productivity and job destruction. Gibbons & Katz (1991) posit that firms lay off their least productive workers first, but that is within a model of employer learning about overpaid lemons, and assumes shock-driven job loss affects all workers equally. The question is considered more thoroughly in the literature on firm turnover and productivity dynamics. The idea is that aggregate productivity changes are driven by reallocation of production from less to more efficient firms. Individual job destruction as it relates to wages is not often considered per se; the connection I make is to infer that firm exit causes job destruction, which pairs with the productivity-wage correlation described above to yield my modeling assumption. See Foster et al. (2008) for a summary of the literature on turnover and productivity. In two examples, Griliches & Regev (1995) show that Israeli firms exhibit lower productivity in the years before they exit,[5] and Aw et al. (2001) find that changes in average industry productivity in Taiwan result from the more productive firms within each cohort being more likely to survive. In the micro search literature, the idea that jobs with some firms are more likely to end than with others has been studied by Pinheiro & Visschers (2015) and Jarosch (2021). Both conclude that wages and job security are most likely positively correlated, yet both treat layoff risk as an innate characteristic of each firm. Bridging the gap between these lines of inquiry, what I propose is a mechanism to explain why survival rates differ across firms, relating productivity to profitability to the job destruction decision.

My model is a steady state one, though I use it to the simulate differential turnover effects of a stylized downturn. From the displaced worker literature we know that there is cyclical variation in job loss, which

[5]Life cycle firm productivity is outside the scope of this paper. As I assume a "firm" here is one job, to extend in that direction I might think of an establishment as a collection of such jobs, with differential destruction leading to changes in average productivity.

itself is greater for less educated workers (Farber, 2017). I interpret education in this case as a proxy for wage differences, because the gradient in Facts One and Two persists within education groups, though I also show that the model can accommodate the evidence if education indicates ability. The broader study of labor market dynamics and productivity over the business cycle typically falls under the purview of macroeconomics. Robin (2011) shows how aggregate productivity shocks can lead firms to terminate matches endogeneously in a model that captures wage dynamics well but does not consider firm-level shocks and is less successful at capturing the dynamics of job destruction. In a calibrated exercise, Burgess & Turon (2010) allow for idiosyncratic and aggregate productivity shocks to better capture cyclical dynamics, with the tradeoff that the distribution of accepted wages for job-to-job movers must be the same as that for people leaving unemployment. Haltiwanger et al. (2018) document that lower productivity firms are more likely to contract in a recession, while Baily et al. (2001) find that plant level productivity is highly procyclical. They also find that equivalent shocks have a greater magnitude effect on productivity for firms that are more likely to downsize, which supports the approach I take later in modeling a downturn. Supposing productivity and wages are linked, related evidence of this downturn effect includes the finding that wage losses of displaced workers are attributable in part to lost firm wage premiums (Fackler et al., 2021), and that recessions increase earnings inequality because of greater churning at the bottom of the distribution (Heathcote et al., 2020).

When workers are subject to the baseline and cyclical layoff risks detailed above, we might expect them to save for a rainy day. In job search models that address these risks, such as Robin (2011), Pinheiro & Visschers (2015), and Jarosch (2021) the prevailing assumption is that workers cannot save, and are in fact risk-neutral. Part of the reason is that consumption data is rarely available, even though Carroll & Samwick (1998) estimates that precautionary saving against shocks accounts for 32 to 50% of wealth. The few empirical search models that do incorporate saving make the tradeoff of taking firms and layoffs as exogenous (Rendon (2006), Lentz (2009), Lise (2013)), or are calibration exercises (Herkenhoff, 2019), and all treat workers has having identical risk preferences even though empirically these are widely dispersed (Kimball et al., 2008). Here I make a different set of assumptions to get the type of model I need to evaluate the policies of interest, incorporating heterogeneous firms, endogenous job destruction, savings, and heterogeneous worker risk aversion.

In regards to the relevant social welfare policy implications in the literature, there is longstanding evidence in the search literature that unemployment insurance prolongs unemployment spells (see Mortensen (1977) for an early treatment, Acemoglu & Shimer (1999) and Lentz (2009) for models of the optimal unemployment insurance level given its disincentives). Acemoglu & Shimer (2000) claim that the prolonged unemployment spells induced by unemployment insurance raise aggregate welfare because workers wait for more productive jobs, which pay higher wages but are also riskier and would otherwise not be as desirable. This is in contrast with the prediction in my model that more productive jobs pay more and are less risky. Rendon (2006) analyzes a large cash transfer and finds that higher initial wealth leads to longer unemployment durations but higher accepted wages. I am not aware of any search models that incorporate a low earnings subsidy as I do, though the principle is well developed (see Moffitt (2002) for a summary).[6] Evaluations find that the earned income tax credit in the United States has a positive effect on employment, especially for single mothers (Hotz & Scholz, 2007), as does the Working Families' Tax Credit in the United Kingdom (Francesconi & Van der Klaauw, 2007).

While the current paper does touch on a broad range of research fields, as discussed in this section they

[6]The United Kingdom implemented a conditional earnings subsidy called the Universal Credit in 2013 as a replacement for other social welfare programs.

actually share many implications in common. What I do in the next section is show how to combine all of the elements from each of these areas of the literature into a single model that is well suited to policy analysis.

3 Model

In this section, I develop a model that can reproduce the gradient of job destruction across the earnings distribution observed in the data in a way that will later allow me to analyze a downturn and the policies of interest. The central features of the model are that firms subject to transitory shocks differ in total factor productivity and that workers with different risk preferences can save to self-insure against layoff risk. After presenting the model, I discuss the underlying assumptions and how they relate to the literature and the main conclusions of the model. In particular, I show how the central implications are robust to potential sorting of high wage workers into high wage firms due to production complementarities. Finally, I demonstrate how I can use the model to simulate an aggregate downturn that is consistent with the patterns in Figure 2 and other empirical evidence about the disproportionate impact on workers at the bottom of the earnings distribution.

3.1 Notation

Firm total factor productivity is x, drawn from distribution $\Gamma(x)$. Transitory productivity shocks are denoted ε, drawn from distribution $F(\varepsilon)$ with density $f(\varepsilon)$. Wages are functions of x, $w(x)$, as is the job termination productivity shock threshold, $\varepsilon^*(x)$. The capital (firm) share of income is α, and the labor share $1-\alpha$. The probability of a worker staying when receiving an outside offer is $P_s(x)$, and the complement, $1-P_s(x)$, is $\overline{P_s}(x)$. The value to a firm is $\Pi(x,\varepsilon)$, denoted Π_ε. Worker value functions are $U(a)$ when unemployed and $V(a,x)$ when employed, where a denotes assets. The shock arrival rates are λ_0 for an offer when a worker is unemployed, λ_1 for an on-the-job offer, η for shocks to the firm, and δ_0 for exogenous job termination unrelated to productivity. A worker's coefficient of relative risk aversion is γ drawn from distribution $F(\gamma)$, and the flow benefit in unemployment is b. The discount factor and rate are β and ρ, where $\beta = \frac{1}{1+\rho}$.

3.2 Environment

The set-up of the model combines the approach of a standard partial job search model with on-the-job search with the framework of transitory productivity shocks in Mortensen & Pissarides (1994). Suppose that infinitely lived workers operate in a stationary environment in discrete time. Workers can be in one of two states: unemployed (U), earning flow benefit b, or employed (E) earning wage w. They seek to maximize utility over time, where utility is a function of consumption c and takes a constant relative risk aversion (CRRA) form, thus exhibiting a finite, constant inter-temporal elasticity of substitution.

$$u(c_t;\gamma) = \begin{cases} \frac{c_t^{1-\gamma}}{1-\gamma} & \gamma > 0 \\ \log(c_t) & \gamma = 1. \end{cases}$$

Whereas this is typically assumed to be the same for all workers, instead suppose that workers are heterogeneous in these preferences, with $\gamma \sim F(\gamma)$. The primary reason for incorporating finite IES is that it gives workers a reason to save. Previous models that consider heterogeneous job destruction rates (Pinheiro & Visschers (2015), Jarosch (2021)) assume risk neutrality, which is mathematically convenient

yet creates an odd pairing where workers purportedly value income stability over time while having linear preferences, or infinite IES. Allowing for saving not only lets me examine self-insurance and consumption smoothing behavior, but it also lets me look at a policy of large one-off asset transfers, which would be irrelevant if workers were risk neutral because they would consume the entire amount immediately. I model heterogeneity in these preferences for two reasons. First, it aligns the model with the data. As I show in Section 4, workers choose a wide range of savings for a given level of earnings. Different preferences can help explain the conditional dispersion in assets. Second, I need it to account for how people respond differently to offers and expectations about the future. Third, it permits workers to sort into jobs based on job security. However, as I show, under the assumptions of the model no such sorting will occur.[7]

As in Lise (2013), assume that workers match randomly with firms, face a known, stationary wage offer distribution $F(w)$ with wages constant for the duration of a match, and have the ability to save out of income at risk-free interest rate r. Workers face the budget constraint

$$c_t + a_{t+1} = i_t + (1 + r)a_t \tag{1}$$

where $i_t = w_t$ or b depending on employment status and consumption is subject to a lower bound on assets, $\underline{a} = 0$, meaning that workers cannot borrow against future income.[8]

Assume that wage offers map to match productivity x through sharing rule α,[9] such that

$$w(x) = (1 - \alpha)x. \tag{2}$$

This implies the existence of a productivity distribution $\Gamma(x)$ and initial flow profits for a firm matched with a worker of

$$\pi(x) = \alpha x. \tag{3}$$

Through random search, workers meet firms and draw offers with arrival rates λ_0 when unemployed and λ_1 when employed. Current matches are also subject to random exogenous termination shocks, arriving at rate δ_0. As in Mortensen & Pissarides (1994), assume as well that matches are subject to transitory negative productivity shocks $\varepsilon \sim F(\varepsilon)$, arriving at rate η. The difference is that unlike in that model, here there is heterogeneity in initial productivity x. Because employment constitutes a fixed-wage contract, following the arrival of shock ε flow profits to the firm become

$$\pi(x; \varepsilon) = \alpha x - \varepsilon. \tag{4}$$

[7] It could, though, potentially occur in a multi-sector form of the model where the sectors differ in terms of exposure to shocks or the worker share of output.

[8] This feature helps fit the empirical fact that many low-income workers live hand-to-mouth, and also motivates precautionary saving in the model. Several search models that allow saving also allow for borrowing (Rendon (2006), Lise (2013)), and Herkenhoff (2019) argue that credit access is important following job loss, yet Rendon (2006) concludes that borrowing constraints are tight and workers can only borrow up to 14% of the present discounted value of income, and the model in Lise (2013) implies much more borrowing than appears in the data. Sullivan (2008) shows that the least well off households do not borrow to smooth consumption in unemployment. While borrowing behavior is outside the scope of this study, if it were included it should tend to reduce precautionary saving yet increase consumption in unemployment.

[9] Determined by Human Resources, or the marginal products of labor and capital in a Cobb-Douglas function with $K, L = 1$.

3.3 Value functions

A firm matched with a worker seeks to maximize the present discounted expected value for the match of

$$\Pi_\varepsilon = \frac{1}{1+\rho}\left(\alpha x - \varepsilon + \lambda_1 P_s(x)\Pi_\varepsilon + \eta \int_0^{\varepsilon^*(x)} \Pi_{\varepsilon'} f(\varepsilon') + (1-\lambda_1-\eta-\delta_0)\Pi_\varepsilon + \delta_0 0\right) \tag{5}$$

by choosing whether or not to continue the match, where P_s is the probability that a worker stays after receiving an outside offer, the discount factor $\beta = \frac{1}{1+\rho}$, $f(\varepsilon)$ is the density of ε shocks, and $\varepsilon^*(x)$ is the threshold for a shock that is large enough for the firm to terminate the match. In a new match, Π_0, $\varepsilon = 0$.

This can be rewritten (see Appendix E) as

$$D\Pi_\varepsilon = \alpha x - \varepsilon + \frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon')d\varepsilon'. \tag{6}$$

where

$$D = \rho + \lambda_1 \overline{P_s}(x) + \eta + \delta_0. \tag{7}$$

Define $\Pi_{\varepsilon^*(x)}$ to be the value where the firm is indifferent between continuing the match or ending it. We know $\Pi_{\varepsilon^*(x)} = 0$ because the firm's outside option is zero,[10] hence the termination threshold $\varepsilon^*(x)$ is

$$\varepsilon^*(x) = \alpha x + \frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon. \tag{8}$$

To interpret that threshold, a firm is willing to keep a match going beyond the point where the flow profit is zero because of the option value associated with drawing a less negative shock in the future. From the worker's perspective, the probability that the firm terminates the match is

$$\delta(x) = \delta_0 + \eta(1 - F(\varepsilon^*(x))). \tag{9}$$

Therefore, a worker meeting a firm with a posting initially producing x receives the offer

$$\{w(x), \delta(x)\} = \{(1-\alpha)x, \delta_0 + \eta(1 - F(\varepsilon^*(x))\}. \tag{10}$$

Proposition 1. *Layoff risk is decreasing in earnings.*

Given the assumptions of the model, wage offers are increasing in firm productivity and layoff risk is decreasing. To see this, take the first derivative of the firm's threshold decision and recall that wages, by definition, are increasing in x (see Appendix E).

Proposition 2. *Workers will not sort into jobs according to their risk preferences.*

Because layoff risk is decreasing in earnings, all workers desire the same types of jobs: a worker who cares more about income wants a high x firm, as does one who cares more about stability. The result simplifies the probability that a worker stays after receiving an outside offer, $P_s(x)$, as values are increasing with

[10] Appealing to free entry, assuming x is drawn after the firm and worker meet.

productivity for everyone. Therefore,

$$\begin{aligned} & P_s(x) = \Gamma(x) \\ \implies & \overline{P_s}(x) = 1 - \Gamma(x). \end{aligned} \tag{11}$$

Workers seek to maximize the present discounted value of expected lifetime utility by choosing consumption and deciding whether or not to accept job offers they receive. The value to a worker depends on her employment status and choice of consumption:

$$U(a) = \max_c \left\{ u(c) + \beta \left(\lambda_0 \int \max \{V(a', x), U(a')\} \, d\Gamma(x) + (1 - \lambda_0) U(a') \right) \right\} \tag{12}$$

$$V(a, x) = \max_c \left\{ u(c) + \beta \left(\lambda_1 \int \max \{V(a', \tilde{x}), V(a', x)\} \, d\Gamma(\tilde{x}) + \delta(x) U(a') + (1 - \lambda_1 - \delta(x)) V(a', x) \right) \right\} \tag{13}$$

where consumption is subject to the budget constraint (1) and the no borrowing constraint $a \geq 0$.

Proposition 3. *Workers with a finite inter-temporal elasticity of substitution will choose to save when employed ($a^*_{t+1} > 0$).*

The reason workers will save is that there is a possibility they will lose their jobs involuntarily. This can happen for two reasons: first, there is a baseline rate of job destruction δ_0 that is unrelated to firm productivity; second, the firm can be hit by a large enough negative shock that swamps the option value of keeping the job filled. Given concave preferences, a worker will want to reduce the variance of consumption across employment states.

Proposition 4. *Given earnings, desired savings are increasing in the level of relative risk aversion.*

The stronger one's distaste for variance in consumption, the stronger one's motivation to reduce that variance. Optimal savings will be where the marginal utility of consumption equals the marginal value of wealth, and the curvature to the marginal utility of consumption is increasing in risk aversion.

Proposition 5. *Given the level of relative risk aversion, desired savings are increasing in earnings.*

This result derives from Proposition 2 in Lise (2013), which also concludes that there is a target level of savings at each level of earnings (proved in Carroll et al. (2003)). Intuitively, the higher one's earnings, the larger the drop in consumption one can expect from job loss. Consider the extreme case where earnings equal the flow benefit from unemployment. In that case, job loss will have no effect on consumption even if savings are zero.

Proposition 6. *Given the level of savings, the reservation wage is decreasing in relative risk aversion.*

This is a standard results, proved in Pissarides (1974) and shown empirically in Feinberg (1977) and Pannenberg (2010). Workers with stronger preferences for smooth consumption will prefer lower, positive income today over no income today and potentially higher income tomorrow.

Proposition 7. *Given the level of relative risk aversion, the reservation wage is decreasing in savings.*

This is a standard result, shown in Danforth (1979) and used in Rendon (2006). Intuitively, when job offers are more frequent for unemployed workers, a wealthier worker can afford to be picky, rejecting offers that a worker with no savings would have to accept while consuming out of savings.[11]

Note that the net effect of these propositions is potential ambiguity in the empirical relationship between risk aversion and reservation wages. More risk averse workers will accept lower offers, yet more risk averse workers will save more, and workers with more savings will wait for higher offers.

Proposition 8. *Unemployed workers do not value layoff risk when evaluating job offers (the reservation wage is independent of risk).*

This result is shown in Pinheiro & Visschers (2015) for the risk neutral case. Observe that a worker is indifferent between two offers when the values are the same:

$$V(a, x') = V(a, x) \implies u(c(a, w')) = u(c(a, w)) + (\delta' - \delta)(V(a, w) - U(a)) \tag{14}$$

and the difference in risk term drops out when the current state is unemployment. Given assets and risk preferences, the reservation wage is independent of the risk offer (see Appendix E). This creates a tension regarding risk preferences, as more risk averse workers have lower reservation wages, which come with a higher probability of sending them back into unemployment.

3.4 Solution

An solution to the model is a set of value functions $U^*(a), V^*(a, x)$ and $\Pi(x, \varepsilon)$ and policy functions for employed and unemployed workers $\phi_u^*(a_t, i_t, s_t)$ and $\phi_e^*(a_t, i_t, s_t)$ such that the worker value functions solve the worker's problem, the firm value function solves the firm's problem, and markets clear. The bond market clears outside of the model in this case, and the goods market clears through the budget constraint for workers, assuming that in addition to the working class there is a separate capitalist class that consumes all firm profits in each period.

3.5 Illustration of model dynamics

To illustrate the offer curves resulting from the model, in this section I calibrate a version of equation 8, parameterizing the firm productivity and shock distributions as described later in Section 5 and using the data from Monte Carlo simulations I use to validate the estimator in Section C.

In accordance with the motivating pattern in Figure 1, Figure 3 simulated from the model implies that the lowest paying jobs will be associated with the highest termination risk, risk will decrease as the wage increases, and ultimately it will level off as productivity outstrips even the largest possible transitory shocks. The observed monthly transition rates into unemployment for the workers simulated from the model shown in Figure 4 decrease in earnings as in the CPS data.

[11] This is not the case in Lise (2013) because differences in search effort across asset levels lead all workers to have the same reservation wage. Although mathematically convenient, this implies that no offers are ever rejected.

Figure 3: Wage-risk offer pairs (calibrated)

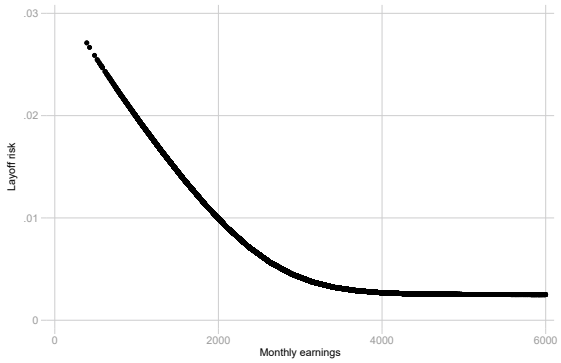

Figure 4: Monthly transition rate to unemployment by income (calibrated)

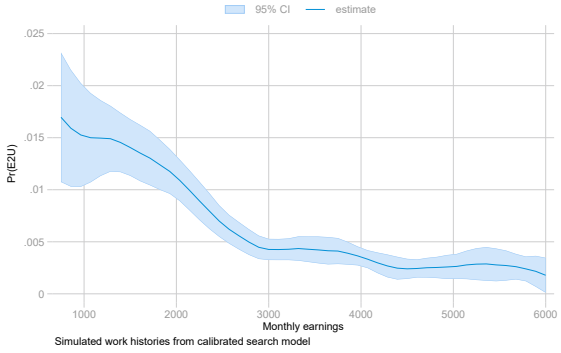

To illustrate the propositions relating to the degree of risk aversion and the level of savings, I trace out the reservation wage curves from the calibrated model in Figure 5. With each risk preference group plotted, with coefficients of relative risk aversion ranging from 0.3 to 1.1, the reservation wage is increasing in savings. Across preference groups, conditional on the savings level the reservation wage is high for the less risk averse group. The curves also illustrate the counteracting effects of risk aversion and savings holdings, which can be seen by imagining a horizontal line drawn at a given reservation wage. For example, in the calibrated model a worker with a risk aversion coefficient of 0.8 and no savings would have the same reservation wage as a worker with a risk aversion coefficient of 1.1 and about $1,500 in savings.

Figure 5: Reservation wage by savings level and degree of risk aversion (calibrated)

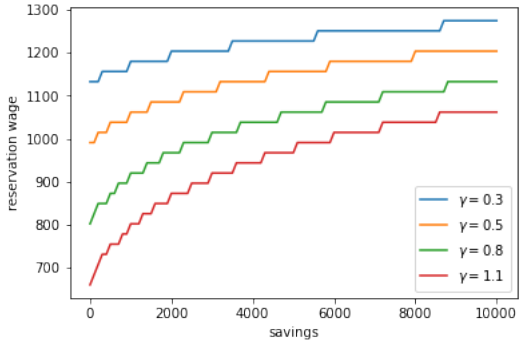

3.5.1 Importance of heterogeneity in inter-temporal preferences

Given these dynamics, allowing for heterogeneity in inter-temporal preferences generates much more dispersion in savings behavior. To see this, I generate two versions of the Monte Carlo data: one with a full distribution of preferences and one where preferences are homogeneous and equal to the mean, with a coefficient of relative risk aversion of about 0.8. Table 1 show how much it matters. Although both assumptions produce similar mean levels of savings, the model with heterogeneous preferences generates substantially wider dispersion. Even with a low degree of relative risk aversion, the 10th percentile of savings is over $5000, vs. $100 assuming a distribution of preferences. The 90/10 ratio is 7.6 with homogeneous preferences compared to 702 with the same mean preference but a continuous distribution around it.

Table 1: Distributions of savings with homogeneous vs. heterogeneous inter-temporal preferences

			Percentile					
	Mean	Std Dev	10%	25%	50%	75%	90%	99%
One preference	24037	13997	5406	12484	24372	33220	41155	61156
Many preferences	27736	39617	100	3149	12940	37774	70260	184810

To provide context for the comparison, I compare the simulated savings distributions to actual data on the distribution of liquid savings for U.S. households in 2016 from the Federal Reserve's Survey of Consumer Finances, summarized in Bhutta et al. (2018). As shown in Table 2, with homogeneous preferences it is difficult to reproduce the fact that about 25% of U.S. households do not have enough savings to cover an unexpected $400 expense, as 99% of households do so in the simulation versus 75% of those in the heterogeneous preferences scenario. Similarly, looking at the rules of thumb for holding savings equivalent to three or six months' worth of expenses, we see that even with fairly a low relative risk aversion level of 0.8, more than twice as many households with homogeneous preferences (88% and 61%) will save at these thresholds compared to the data (40% and 28%), whereas heterogeneous preferences with the same mean can get much closer (58% and 42%). In the right tail, again a model with many preference levels is better able to capture the 20% of actual households with at least a year's worth of expenses in liquid savings than a model of common preferences, where almost no one chooses to save that much.

Table 2: Comparison of simulated savings distributions to U.S. household savings

	% households with savings for the following expenses				
	$400	1 month	3 months	6 months	12 months
One preference	99%	97%	88%	61%	1%
Many preferences	80%	75%	58%	42%	18%
Actual U.S. data	76%	60%	40%	28%	20%

Source: Bhutta et al. (2018)

3.6 Discussion of main assumptions

In this section I discuss the main assumptions I make to construct the model. I explain what I gain from each and what evidence there is in favor or against it in the literature. While some of the assumptions are strong relative to what is typical in a structural job search model, my emphasis here is on the dynamics of job destruction and labor market policy responses to a downturn, and these assumptions let me relax other assumptions so that I can match a specific empirical pattern that those models cannot.

3.6.1 Wage determination process

I assume that the wage offered is a direct function of firm total productivity and does not result from bargaining over a surplus. Although bargaining is more standard in a search model, survey evidence reports that two-thirds of workers did not negotiate their current salary (Hall & Krueger, 2012), and empirical estimates of wage bargaining shows that it is uncommon at low and medium wage jobs (Cahuc et al. (2006), Card et al. (2016)). In this paper I am particularly interested in how lower wage workers are affected by job destruction dynamics, and these are the workers for whom there is the least evidence of bargaining. For workers who do bargain, estimated magnitudes of the bargaining are very small, with an elasticity with respect to increases in the outside option on the order of just 0.1-0.2% percent (Caldwell & Harmon (2019), Lachowska et al. (2021)).

The conclusion about the relationship between wages and layoff risk does not necessarily require as strong an assumption as exogenous wage determination. The conclusion should hold so long as 1) wage offers are increasing in match productivity, an empirical regularity reviewed in Card et al. (2018), and 2) wages do not rising significantly faster than productivity to the point where more productive firms earn lower profits. Estimates of the elasticity of wages with respect to productivity tend to be around 0.05 to 0.10 (Card et al., 2018), which is far below that threshold.

One reason I believe bargaining is not crucial here is that the Mortensen & Pissarides (1994) model with bargaining still results in firms having a productivity threshold below which they will terminate a match, and this threshold is increasing with the initial productivity of the match. That model delivers a tractable bargaining solution by assuming no on-the-job search and homogeneous, risk-neutral workers. Once workers are allowed to differ and to weigh outside offers, which are valued differently over time according to worker preferences and the current wage, the bargaining solution quickly becomes unwieldy. To flip the implications in the model and get layoff risk to increase with wages, bargaining power would need to rise so steeply with match productivity as to cause profits to decrease, which does not accord with intuition, the magnitude of the bargaining power estimates in the literature, or the empirical evidence that, if anything, the labor share of income is decreasing with firm productivity (Dorn et al., 2017). Because the bargaining approach does not deliver a fundamentally different conclusion regarding the relationship between productivity and job security, while adding firm and worker heterogeneity and on-the-job search to the model, I make the simplifying

assumption that there is no bargaining. In Appendix A I show how bargaining could be incorporated in a simplified version of the model without changing the relationship between wages and layoff risk.

3.6.2 No wage renegotiation

I assume that wages are fixed for the duration of a match. That is, firms have a set pay for a position and do not like to change it.[12] This is equivalent to supposing that workers, who are risk averse, are fully insured against transitory shocks by the risk-neutral firms, after Azariadis (1975). Recent evidence that firms act as insurers comes from Guiso et al. (2005), who find that transitory shocks to firms are not generally passed through to workers; Altonji & Devereux (1999), who observe that downward wage cuts without changes in job status are rare; Haefke et al. (2013), who show that new hires have much more flexible wages than current workers; and Barattieri et al. (2014), who conclude that workers who switch jobs have a much higher probability of changing wages than workers who continue to be employed at the same firm. Jardim et al. (2019) and Elsby & Solon (2019) argue that nominal wage cuts are more prevalent than found in studies using household survey data, finding that about 15-25% of job stayers experience nominal wage cuts in a year, but this is contradicted by Grigsby et al. (2021), who use administrative data to show that only 2.5% of job stayers experienced a wage cut between 2008-2016. Lamadon et al. (2019) show that pass-through of shocks, though low, is not zero. Kaur (2019) extends the analysis to India, finding that positive transitory shocks may increase nominal wages but negative shocks do not.[13] Fallick et al. (2020) shows that most U.S. employers did not adjust wages during the Great Recession even though unemployment spiked.

Qualitatively, there is a long documented history that firms do not like to adjust wages downward, and the current push for wage transparency is specifically intended to prevent firms from paying different wages to similar workers. Bewley (2009) finds that employers prefer not to cut wages because they believe it will reduce worker morale and engagement with the firm's objectives. Similarly, Campbell III & Kamlani (1997) documents resistance to wage reductions particularly because of employee morale and effort issues and the risk of losing their most productive workers. One implication of wage rigidity is that unemployed workers could underbid current employees by being willing to accept a lower wage for an equally productive match. Solow (1990) argues that the absence of underbidding could explain observed patterns in downward wage rigidity, while Fehr & Falk (1999) finds it is actually firms that refuse to accept that underbidding. The explanation offered is that since effort is often unobservable, firms prioritize fairness to maintain worker cooperation, which wage reductions put in jeopardy. This is consistent with survey results in Kaur (2019) showing workers see wage cuts as unfair, and that cuts reduce effort.

The alternative in the other direction is to allow for full pass-through of shocks as in Mortensen & Pissarides (1994). That assumption has been criticized in the macro literature for being unable to capture the volatility of unemployment fluctuations (Shimer, 2005). In what direction would renegotiation changes the results? First, in the data I would expect frequent wage adjustment within job spells. Second, in the model firms should keep jobs open for longer because they could endure larger shocks. However, as I show in Appendix A, the order ranking between firm productivity and job termination probability would be unaffected. Although a model of imperfect pass-through is the most realistic, I opt for the stronger assumption because it is simpler, because I interpret the empirical evidence to lean towards wages not adjusting frequently, and because incorporating updating would not affect the job destruction dynamics I am interested in capturing.

[12] An assumption more readily accepted by anyone who has had to deal with the Human Resources Department.

[13] Inflation may allow firms to cut real wages without adjusting nominal wages. I abstract away from that here, as it affects all firms equally and is not a viable individual firm response to a transitory shock.

3.6.3 Homogeneous worker skill

I assume that workers do not have innate productivity differences. While it is not uncommon in the search literature, clearly this abstracts from reality. Here I discuss that assumption and show how skill differences could enter into the model as a production complementarity, motivated by the finding in Lamadon et al. (2019). A first reason for maintaining the assumption is comparability: the most closely related work on savings and search (Rendon (2006), Lise (2013)), and on heterogeneous job risk and search (Jarosch (2021), Pinheiro & Visschers (2015)) make the same assumption. Shutting down the dimension of innate worker skill differences allows for more direct examination of the dynamics of layoff risk itself. That being said, while the literature broadly finds that "trends in aggregate wage dispersion closely track trends in the dispersion of productivity" (Card et al., 2018), it may be because high skill workers are sorting into high productivity firms (Winter-Ebmer (1995), Abowd et al. (1999)). As discussed in Postel-Vinay & Robin (2002), reaching such a conclusion by examining worker fixed effects in labor market histories may suffer from an endogeneity problem if those prior trajectories result from search frictions and not innate skill differences. That is, workers who initially match randomly with a lower productivity firm could churn into unemployment more frequently over their careers than identical workers who initially match with a higher productivity firm.

The current paper is agnostic on the question of worker sorting because any such behavior should not affect the fundamental relationship between earnings and layoff risk. Although the prevailing earnings distribution would differ by skill in the model, layoff risk conditional on earnings would be unaffected. To see how, suppose that skill s enters into the model by multiplying firm productivity:

$$x' = sx \tag{15}$$

such that total productivity is the product of the firm component and a worker component. Wage offers will reflect this total productivity:

$$w(x; s) = (1 - \alpha)x'. \tag{16}$$

The firm's threshold for terminating a match will also be a function of the combined product, but that is equivalent to the problem facing a firm in the baseline model that has productivity $x = x'/s$.

To demonstrate this result, I simulate a calibrated version of the model where half of the workers are "low skill" (s = 1) and half are "high skill" (s = 1.2). Figure 6 shows how layoff risk relates to earnings. Conditional on earnings, both skill groups experience the same risk, even though as shown in Figure 7 the high skill group has higher earnings. Across and within skill groups, workers in more productive matches will still earn more and enjoy greater job security.

Figure 6: Quarterly transitions into unemployment by earnings and skill (calibrated)

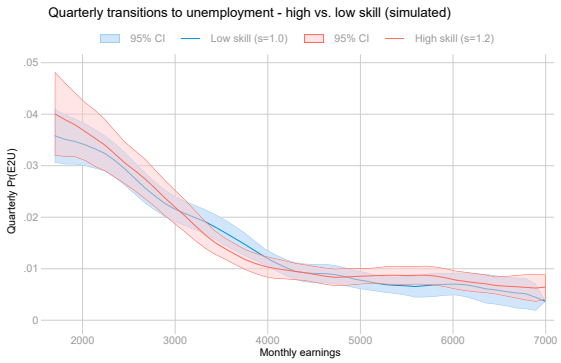

Figure 7: Distribution of earnings by skill (calibrated)

One could more directly align the model with an AKM approach by supposing that firm and worker effects enter additively, an approach which has been shown to explain a substantial amount of the residual variation in wages (Card et al., 2016). If the worker effect is s, let match total factor productivity be

$$x' = x + s \tag{17}$$

with corresponding wage offer

$$w(x; s) = (1 - \alpha)x'. \tag{18}$$

As in the multiplicative case, the inverse wage-layoff risk relationship follows in the same way as in the base case. Higher worker effect workers will receive higher offers, yet layoff risk conditional on total match productivity will be unchanged. Note that this suggests a mechanism for partial worker sorting. If higher skill workers mechanically are faced with higher productivity matches, those matches will be less likely to terminate, giving the higher skill workers more time and opportunity to climb the job ladder. Sorting is only partial because labor market frictions prevent all workers from racing to the highest paying firms right away.

If instead innate worker differences affect the labor share of income, a mechanical tradeoff will arise. The workers who earn a higher fraction of the marginal product will face greater termination risk at all wage levels, as the firm will have a lower profit margin given the match productivity. Note that this implies lower wages are not a pure lose-lose for the "underpaid" workers: lower pay conditional on productivity

mechanically comes with higher job security.[14]

3.7 Incorporating a downturn

By combining search frictions and the ability to save, the model provides a novel framework for analyzing the level and distribution of consumption in an aggregate downturn and the important role of prior wealth accumulation in moderating its effects. As suggested by the trends during the Great Recession and the Covid-19 pandemic, the impact is likely unequal because layoff risk appears to increase relatively more for lower earners. In this section I demonstrate how to use the model to simulate an unexpected aggregate downturn.

To produce an aggregate downturn in the model, I suppose that there is a permanent, unexpected shock that drives down all match output by s.[15] Technically, this is not exactly a recession, as this is a steady state model while recessions are by nature cyclical. We can either view this type of downturn as more like the permanent local labor market shocks experienced by many U.S. commuting zones exposed to Chinese export competition in the early 2000s (Autor et al., 2016), or simply as a stylized depiction of how the labor market responds at the onset of a recession before any recovery begins.[16] Furthermore, this stylized consideration of an aggregate downturn can capture several findings in the macroeconomic literature about productivity and cyclical labor market sorting. In particular, I will show that lower productivity firms are more likely to contract in a recession (Haltiwanger et al., 2018). The model also accommodates the evidence that wage losses of displaced workers are attributable in part to lost firm wage premiums (Fackler et al., 2021), and that recessions widen earnings inequality because of increased churning at the bottom of the distribution (Heathcote et al., 2020).

Suppose that a match previously producing x now produces $\tilde{x} = x - s$. I make this additive assumption because it aligns with the finding by Baily et al. (2001) that the fall in productivity from a negative shock for firms that tend to downsize is significantly larger in magnitude than that for other firms facing the same degree shock. Because prevailing wages are rigid, what previously earned the firm a net flow profit of $\alpha x - \varepsilon$ now earns $\alpha x - s - \varepsilon$. Profit margins have narrowed, meaning a smaller transitory shock could lead to job destruction and hence layoff risk is higher. This can be seen in the new equation for the firm's termination threshold:

$$\tilde{\varepsilon}^* = \alpha x - s + \frac{\eta}{\tilde{D}} \int_0^{\tilde{\varepsilon}^*} F(\varepsilon) d\varepsilon \tag{19}$$

For newly formed matches, assume that offers adjust to the downturn conditions. The new distribution of wage and corresponding layoff risk offers maps from $\tilde{x}$, e.g. $w(\tilde{x}) = (1-\alpha)(\tilde{x})$. This is an order preserving shift in the values offered to unemployed workers by these new firms. For firms in current matches at the onset of the downturn, the probability that their worker leaves for another firm has become slightly more complicated, as outside offers reflect the new distribution whereas the current offer has a higher wage and higher risk than it would if the shock were passed through to the worker. In practice, this nuance should not affect layoff risk much because it only enters into the termination threshold as a multiplier on the option

[14] While the notion that workers with greater bargaining power take on high layoff risk appears odd in the context of skill, it is more plausible for a discrete difference such as gender, where differences in bargaining power have been detected (Card et al., 2016). Exploration of differential bargaining or sector sorting is outside the scope of this paper.

[15] While combinations of shifts in multiple parameters could affect layoffs, such as a change in the shock arrival rate or the variance of those shocks, this is the simplest, most direct approach that reproduces the patterns observed in the data.

[16] A stylized recovery could then be simulated, if desired, by undoing the aggregate shock.

value of continuing the match. Recalling that $D = \rho + \lambda_1 \overline{P}_s(x) + \eta$, and probability is bounded between 0 and 1, the unconditional bounds on η/D are $\left\{\frac{\eta}{\rho+\lambda_1+\eta}, \frac{\eta}{\rho+\eta}\right\}$. Since outside offers have in general become worse, the probability of leaving will decrease, raising the lower bound to $\frac{\eta}{\rho+\lambda_1\overline{P}_s(x)+\eta}$. In the policy section simulations, I leave the probability unchanged because accounting for this nuance would complicate the code considerably without materially affecting the results.

Having developed a model with the necessary properties to capture the mechanics of differential job destruction across the earnings distribution in this section, in the next section I turn to describing the data that will let me estimate the structural parameters of the model.

4 Data

In this section, I describe the main data source used to estimate the model and details the construction of the sample. The source is the Panel Study of Income Dynamics (PSID), which is a national longitudinal household survey based out of the University of Michigan that has been interviewing the original sample members and their descendants since 1968.[17] Since 1999, the PSID has collected detailed information about household wealth and consumption in each biennial survey wave. Consumption data is rare, and the fact that it collects such data in conjunction with the standard variables makes the PSID one of if not the only longitudinal panels in the United States that reports earnings, employment status, savings, and consumption. Saving stocks and annual consumption are reported once per wave. Beginning with the 2003 wave, the PSID has constructed a monthly employment matrix for primary respondents covering up to four distinct jobs in each survey year.[18] Together, these features of the data provide the necessary inputs for estimating the model developed in the paper: consumption, savings, monthly employment status, and earnings. I use the full set of survey waves containing this information, which are those from 2003 to 2019.

Starting from the full set of primary respondents in the 2003 to 2019 waves, I limit the sample to individuals ages 21-64 who report that they are either working or looking for work, removing respondents who are self-employed, retired, disabled, students, or in an institution. Second, I remove female household heads because of the PSID sampling procedure. Women can only be considered heads of household if they are unmarried. A married woman's husband becomes the new head, leading to a selection problem. For a similar reason, I remove what are called "non-gened" males. These are men who married into a survey family as opposed to being born into one. The PSID does not follow these men if they get divorced; therefore non-gened males can only be married, creating an inverse problem to that caused by the rules for women as household heads. Finally, I trim assets, consumption, and earnings observations at the 1 and 99 percentile of the data in each wave to minimize the influence of outliers that are likely to be caused by errors in reporting or classification of responses (for instance, if in the data earnings are coded as per hour instead of per week). The resulting sample includes 4284 unique individual respondents with a total of 16505 wave-to-wave observations.

The specific measurements used as model inputs are defined as follows. Employment status is treated as binary for each month. A worker is coded as unemployed in a month if they are unemployed at any point during the month. Reported earnings for each job are converted to monthly values. To convert, I divide annual wages by 12, multiply biweekly wages by 2, multiply weekly wages by 4, multiply daily wages by 20 and multiply hourly wages by 4 times average hours worked per week. I divide annual consumption by

[17] `https://psidonline.isr.umich.edu/`

[18] Wave year t-1. Monthly data for wave year t-2 is available but not in a format I can use to estimate the model without imposing extreme assumptions about earnings and job transitions. Thus, I have a monthly panel in alternating years.

12 to obtain a monthly average. Savings are defined as liquid savings that might be accessible to finance consumption in unemployment, which I calculate as the sum of cash, bonds, stocks, and IRA annuity holdings. All values are converted to 2010 dollars using the U.S. consumer price index. Table 3 displays descriptive statistics for the constructed sample.

Table 3: Descriptive statistics for PSID sample

Waves	8			UE rate	0.041
Workers	4284			U2E monthly	0.110
Worker-waves	16505			E2U monthly	0.005
Worker-months	232119			J2J monthly	0.007
	Mean	Median	Std Dev	Min	Max
Waves per worker	3.85	3.00	2.31	1.00	8.00
Monthly Earnings	4591.38	4257.61	2058.26	1139.88	12747.78
Savings	49443.82	9072.38	129592.00	0.03	959890.30
Savings/earnings	11.34	2.23	33.36	0.00	825.23
Consumption	4022.18	3711.59	1957.26	604.97	14625.92

To adjust for level differences on the basis of observable heterogeneity, I follow Jarosch (2021) and Lise (2013) and residualize certain measures by projecting them onto a set of covariates, removing the predicted variation, and using the remaining variation as the sample measurement. To adjust assets for family composition and demographic background, I project the log of assets + 1 onto marital status, the number of children, race, education, and year, and then translate the distribution by adding back in the mean of log sample assets. For wages, I project onto a Mincerian specification to get residual wages, including time fixed effects because this is a steady state model. To do so, I regress log wages on age, age-squared, education, and year dummies, and restore the mean of log wages to the residual values to generate the residualized wage distribution.

Two key features of the data that the estimation will match are the transition rates into unemployment across the earnings distribution, which will be used to make inferences about layoff risk, and the ratio of savings to earning across the earnings distribution, which will be used to make inferences about risk preferences. Figure 8 shows a local polynomial plot with a 95% confidence interval of the quarterly transition rates into unemployment for the constructed sample. I present quarterly rates because the dynamic moments I match in the estimation will be from quarter to quarter. The figure demonstrates the same pattern observed in the motivating pattern in the CPS data from Figure 1, whereby layoff risk is highest for the lowest earners, slopes down up to around median earnings, and then levels off. This information will be used in the estimation to characterize the distribution and arrival rate of shocks faced by firms.

Figure 8: Quarterly transition rate to unemployment in the PSID sample

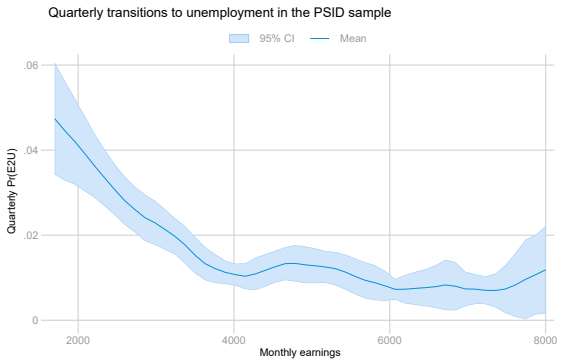

Figure 9 shows a local polynomial plot with a 95% confidence interval of the stock of savings holdings as a function of monthly earnings in the PSID sample. Mean savings rise nearly linearly with earnings at a ratio in the neighborhood of 10 to 1. This information will be used in the estimation to locate the mean of risk preferences among workers. Within each earnings band, there is a fairly wide distribution of savings to earnings ratios, depicted with kernel density plots in Figure 10. The distribution has a higher central tendency and wider dispersion in higher earnings buckets relative to lower ones. This information will be used in the estimation to characterize the variance of risk preferences among workers.

Figure 9: Savings per earnings in the PSID sample

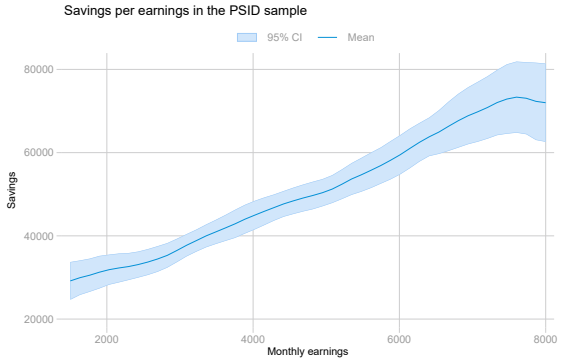

Figure 10: Distribution of savings to earnings ratios in the PSID sample

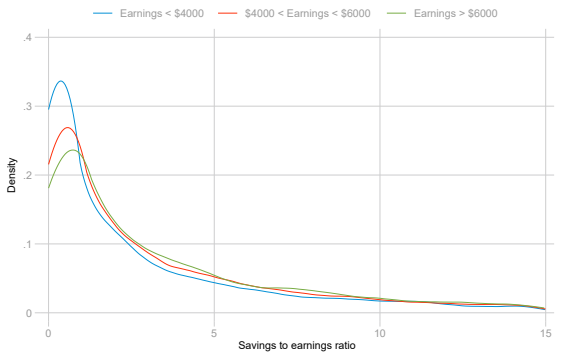

In building the model I will allow for workers to save. Figure 11 presents descriptive evidence that savings is important in order for the model to fit the data. It plots a local polynomial regression of mean monthly consumption against monthly earnings for the PSID sample. Most workers consume less than they earn, and the consumption rate decreases with earnings. Mean consumption does exceed earnings for the very lowest earners. Borrowing, which I do not model, could be one explanation, but two others consistent with my no-borrowing assumption are that 1) the PSID consumption data is reported net of any government transfers, 2) some of these workers could be consumption smoothing by dissaving.[19]

Figure 11: Consumption as a function of earnings in the PSID sample

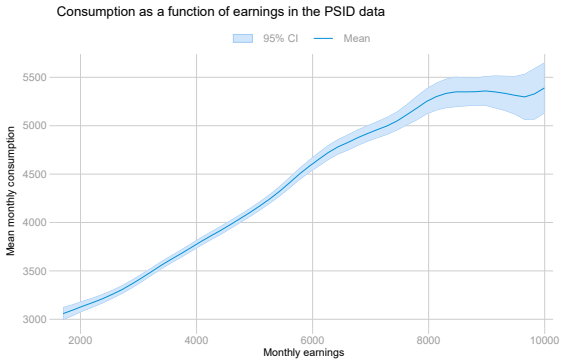

In this section I introduced the PSID panel and the construction of the estimation sample. I presented a descriptive overview of the sample and highlighted the key elements of the data that will inform the estimation process. The next section explains the estimation approach and how each element of the data maps to the structural parameters I want to measure.

5 Estimating the model

In this section, I describe how I estimate the structural parameters of the model using the PSID data introduced in the previous section. First, I review the set of parameters and the additional parametric assumptions I impose in order to identify them empirically. Then, I discuss the set of moments I match in the simulated method of moments estimation and how they map from the data to the parameters. Finally, I discuss how restrictive the identifying assumptions I make are, and how dependent I am on those assumptions in order to estimate the model.

The full set of parameters in the model is the following:

$$\Theta = \{\lambda_0, \lambda_1, \delta_0, \eta, \alpha, b, r, \rho, F(\gamma), \Gamma(x), F(\varepsilon)\}\,. \tag{20}$$

To recap from the model section, these are, in order, the arrival rate of job offers in unemployment and employment, the baseline exogenous job termination rate, the arrival rate of shocks to the firm, the worker share of output, the flow benefit in unemployment, the risk-free interest rate, the discount rate, the distribution of risk preferences, the distribution of firm total factor productivity, and the distribution of firm shock magnitudes.

[19] A third explanation outside the scope of the model is that this is household consumption plotted against individual labor earnings. Some households may have additional sources of income.

I fit most of the parameters by the simulated method of moments, but first I calibrate a few and estimate one other in a first stage. Following standard practice in the literature, I calibrate the interest rate and discount rate: $r = 0.00125$ and $\rho = 0.006$ ($\beta = 0.996$). I also fix the worker share of income as $1 - \alpha = 0.6$, to match the aggregate labor share of income prevailing in the United States between 2002-2016 (University of Groningen and University of California, Davis, 2021).[20] Next, extending an argument from Flinn & Heckman (1982) about identifying the reservation wage by the minimum wage in the sample, I estimate the flow value of unemployment b as minimum value of consumption in the sample: $\widehat{b} = \min \{c_i\}_{i=1}^{N}$.[21]

Simulated moments (McFadden, 1989) works by calculating that set of moments using the sample data (m_d), simulating the same set of moments from the economic model using a candidate set of parameters (m_s), and computing the distance between the two vectors. The objective is to choose the set of parameters that minimizes this weighted distance:

$$\widehat{\Theta} = \arg\min_{\Theta} (m_d - m_s)' W (m_d - m_s) \tag{21}$$

where to make the problem scale well to ease estimation I follow a typical practice of setting W as a diagonal matrix containing the inverse moment variances, which, also following convention, I obtain by bootstrapping the sample.

To fit the remaining parameters, I need to choose an appropriate set of moments m_d in the data that should uniquely or jointly map to each parameter. The full set of moments I use in estimation is listed in Table 4. For some parameters, like the job offer arrival rates, there are standard moments that many papers use. For instance, the unemployment rate, the quarterly transition rate from unemployment to employment, and the rate of wage increases target the offer arrival rates λ_0 and λ_1. To target the wage offer distribution, I choose several moments from the full wage distribution among employed workers and also just for wages accepted out of unemployment, specifically the mean, median, and standard deviation, as well as the ratios of the 90th to 50th and 75th to 25th percentiles of accepted wages. In many papers, given the model assumptions the distribution of wages accepted out of unemployment exactly identifies the wage offer distribution, but that is not the case here because workers all have unique reservation wages depending on their preferences and savings levels. For the wage distribution itself, I assume a flexible parametric distribution for log firm productivity x, supposing it is distributed Beta(α_x, β_x). This maps to wages through α. I then need a way to scale it appropriately, as a beta distribution has support between 0 and 1. To do so, I follow Bontemps et al. (2000) who show that the upper and lower bounds for the support of the wage distribution, $\underline{w}$ and $\overline{w}$ can be estimated consistently by the sample minimum and maximum.

The two less standard types of parameters I fit here are those related to firm shocks and endogenous job destruction and the distribution of worker risk preferences, so I will discuss the choice of those moments in more detail. First, the parameters relating to layoff risk. For these I choose moments capturing a range of conditional transitions into unemployment. Recall that the overall termination rate in the model is $\tilde{\delta}(x) = \delta_0 + \eta(1 - F(\varepsilon^*(x)))$. For the highest paying firms $F(\varepsilon^*(x)) \to 1$, meaning moments for transitions to unemployment among high earners identify δ_0. Likewise, for the lowest paying firms $F(\varepsilon^*(x)) \to 0$, hence I choose the transition rate into unemployment for the lowest earners to match the arrival rate of productivity shocks η. Lastly, I make a parametric form assumption for $F(\varepsilon)$, assuming it follows an exponential

[20] Recognizing that there is a strand of the literature that seeks to explain why this share appears to be falling more recently (Grossman & Oberfield, 2021).

[21] Relying on this order statistic for identification may not be valid if the model allowed for borrowing, and the statistic itself could be sensitive to sample trimming. I check this second possibility in the sensitivity analysis.

distribution with parameter θ_ε. I target θ_ε by the overall average transition rate into unemployment and the rates of intermediate earners across the distribution.

To find moments to match for the risk preference parameters, consider that under the buffer stock theory of precautionary saving (Carroll & Samwick, 1998), and as in Lise (2013), workers will have a target savings-to-earnings ratio given their level of earnings and the risk of unemployment. More risk averse workers will choose a higher target at all levels of income. To demonstrate, in Figure 12 I plot the density and relationship of the savings-to-income ratio for three groups of risk preferences based on a calibrated simulation of the model. The lower the level of risk aversion, the lower the desired level of savings at all wages, and the more mass there is concentrated at low levels of the the savings-to-income ratio. Technically, what we observe is a mixture of such ratios over the people with each risk preference, so a parametric assumption is important to be able to recover the distribution. I assume a fairly flexible parametric form for risk preferences, supposing they follow a beta distribution Beta$(\alpha_\gamma, \beta_\gamma)$, and then I target those parameters with moments capturing the mean and spread of the ratio between savings and earnings for several groups across the earnings distribution.[22]

An alternative approach would be to make use of the stated risk preferences of some of the workers. These are available for household heads who were interviewed in the 1996 wave, and are a popular source of variation in papers that study risk preferences. However, as I show in Appendix D, these responses do not correlate with savings behavior as one would expect if they represented preferences over income risk. Therefore, I prefer to model preferences as unobserved heterogeneity.

Figure 12: Empirical identification of risk preferences

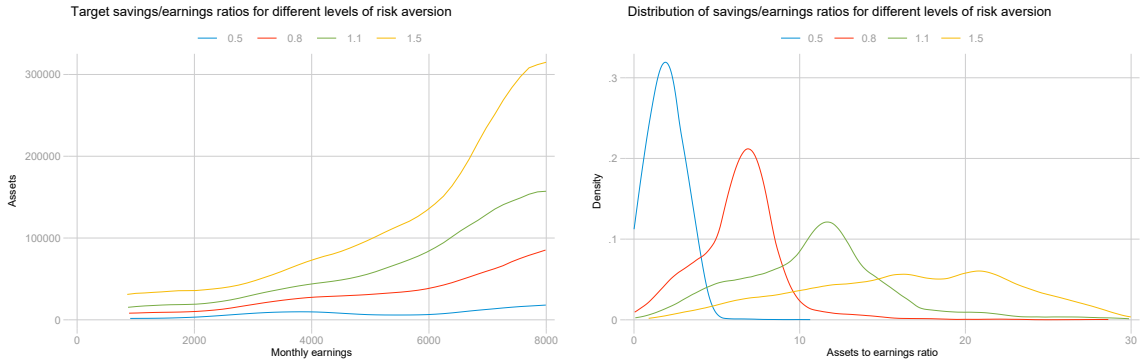

[22]Since support for the beta distribution is between 0-1, I scale those values to 0-3. This is conservative given the claim in Chetty (2006) that the theoretical upper bound on the coefficient of relative risk aversion is about 2, yet restrictive given the much higher estimates of 5 or more in Kimball et al. (2008). I ultimately chose the scale I did after estimates with a much wider scale consistently supported Chetty (2006), yielding a distribution with all of the mass below 3.

Table 4: List of sample moments targeted in estimation

1	UE %	13	Median wage	25	90:50 quantile ratio of wage U->E
2	U->E %	14	Median wage U->E	26	75:25 quantile ratio of wage U->E
3	E->U %	15	90th quant. wage U->E	27	% wage increases
4	E->U wage quant. 1-10	16	Mean ln(savings)	28	Mean assets b/w wage quant. 75-99
5	E->U wage quant. 10-25	17	Std ln(savings)	29	Mean ratio assets:wage wage quant. 20-40
6	E->U wage quant. 25-50	18	Median ln(savings)	30	Mean ratio assets:wage wage quant. 40-80
7	E->U wage quant. 50-75	19	Mean ratio assets:wage	31	Mean ratio assets:wage wage quant. 80-99
8	E->U wage quant. 75-99	20	Std ratio assets:wage	32	Std ratio assets:wage wage quant. 20-40
9	Mean wage	21	Median ratio assets:wage	33	Std ratio assets:wage wage quant. 40-80
10	Std wage	22	Mean assets given wage quant. 10-25	34	Std ratio assets:wage wage quant. 80-99
11	Mean wage U->E	23	Mean assets b/w wage quant. 25-50		
12	Std wage U->E	24	Mean assets b/w wage quant. 50-75		

A single iteration of the estimator consists of finding a fixed point for the value functions given the candidate parameters, simulating the worker histories, calculating the simulated moments, and evaluating the objective function. The estimation procedure involves 20,000 such iterations carried out using the controlled random search with local mutation global optimization algorithm described in Kaelo & Ali (2006) as implemented in the NLopt toolkit for Python (Johnson, 2014). I then polish the global estimates by running the NLopt implementation of a local subplex[23] algorithm (Rowan, 1990) for 1,000 additional iterations.

While the above discussion of how the structural parameters are identified is not treated formally, I do show a form of empirical identification in Appendix C, where I generate a Monte Carlo sample of work histories from the model and then run the estimator on that data. As shown, it successfully recover the true parameters. Having established an estimation strategy and validated it against data generated by a known process in this section, in the next section I run the estimator against the PSID sample and analyze the results.

6 Estimates and model fit

In this section, I present the results from estimating the model against the PSID data. I contextualize the results given previous estimates in the literature and evaluate how well the estimates fit the data. I also present the results of a sensitivity exercise investigating how dependent the estimated values are on the choice of optimization algorithm and some of the parametric assumptions.

Table 5 displays the parameter estimates. To obtain the standard errors given in parentheses below the estimates, I took 1000 bootstrap draws and ran the estimator on each of those samples. The estimates seem reasonable given other findings in the literature. For instance, unemployed workers are estimated to have about a 19% chance of receiving an offer in a given month, and employed workers have a lower arrival rate of about 9%. The on-the-job offer arrival rate is consistent with the calibrated ranges implied for the U.S. labor market by Hornstein et al. (2011), who argue that it takes a rate between about 7-14% to reproduce the job-to-job transitions observed in the data. The arrival rate of offers in unemployment is in a range consistent with the finding in a BLS analysis of a 2018 CPS data supplement that active job seekers have on average a 26.24% chance of receiving a job offer (Dalton & Groen, 2020). This exceeds the monthly transition rate out of unemployment of 8% observed in the data because heterogeneity in worker preferences and savings implies that some offers will be rejected even though no offers will be below the minimum reservation wage. In a model with sequential bargaining, the arrival rate would exactly equal this transition rate, as every

[23] "A variant of Nelder-Mead that uses Nelder-Mead on a sequence of subspaces," (Johnson, 2014).

meeting between a firm and worker forms a match. My results imply that workers accept about one out of every three offers they receive, which rationalizes the low median earnings offer I find of about $25,000 per year relative to the prevailing median earnings level in the U.S. of $30,200 in 2012. The arrival rate estimates mean about that half of unemployed workers will take longer than three months to receive their first offer. Therefore, every time a worker loses their job, search frictions imply a substantial loss of income. This point is highlighted Burdett et al. (2020), who find that an unemployment spell in Germany costs 8–9% of expected lifetime discounted earnings. The gap between the arrival rate and the acceptance rate also has important implications for the earnings subsidy I analyze in the next section. The larger the gap, the more leverage the subsidy has to incentivize workers to accept an offer. If everyone was already accepting every offer, no amount of subsidy could increase the acceptance rate.

Table 5: Parameter estimates

parameter	estimate
α_x	3.2827
	(2.4260)
β_x	9.3154
	(4.8067)
λ_0	0.1861
	(0.0878)
λ_1	0.0863
	(0.0247)
θ (x100)	0.0603
	(0.0187)
δ_0 (x100)	0.1602
	(0.0920)
η	0.0496
	(0.0426)
α_γ	3.0326
	(1.0526)
β_γ	4.7141
	(4.8274)
b	604.9742
	(3.6191)
$\min(x)$	1899.8007
	(3.7141)
$\max(x)$	21246.3033
	(35.8364)

Bootstrap standard errors in parentheses.

Table 6: Comparison of sample and simulated moments

Sample Moments					Simulated Moments				
M01-04	4.18E-02	4.21E-01	1.42E-02	5.15E-02	M01-04	3.54E-02	3.76E-01	1.26E-02	4.69E-02
M05-08	2.12E-02	1.03E-02	7.38E-03	5.67E-03	M05-08	2.42E-02	1.06E-02	3.58E-03	3.41E-03
M09-12	3.36E+03	1.87E+03	2.92E+03	4.27E+03	M09-12	2.43E+03	6.98E+02	2.26E+03	4.17E+03
M13-16	4.61E+03	2.07E+03	5.74E+03	8.68E+00	M13-16	4.54E+03	1.91E+03	3.40E+03	8.78E+00
M17-20	2.74E+00	9.12E+00	1.14E+01	3.37E+01	M17-20	2.56E+00	9.25E+00	1.03E+01	3.02E+01
M21-24	2.23E+00	8.05E+00	8.43E+00	9.03E+00	M21-24	2.51E+00	8.19E+00	8.56E+00	8.88E+00
M25-28	9.47E+00	1.96E+00	2.12E+00	2.09E-02	M25-28	9.79E+00	1.51E+00	1.42E+00	2.30E-02
M29-32	7.53E-01	2.32E+00	5.00E-01	2.51E+00	M29-32	1.44E+00	1.44E+00	4.23E-01	2.49E+00
M33-34	1.75E-01	2.75E+00			M33-34	3.28E-01	2.27E+00		

Table 6 lists the sample value and the simulated value for the full set of moments. As a reminder, explanations for what each moment represents appears in Table 4. Because the emphasis of the model is on differential layoff risk and heterogeneous savings behavior, I am going to focus more on those sets of

parameters in exploring the results checking the fit of the model. The moments related to job destruction are numbers 3 to 8: overall transitions into unemployment and the transition rate for quantile subgroups across the earnings distribution. The estimator fits these fairly well, which I also show in Figure 13, which plots a local polynomial regression for the transition rate into unemployment across the earnings distribution. This pattern of decreasing transitions into unemployment in Figure 13 cannot be matched by a standard search model that has a single exogenous job termination arrival shock. Jarosch (2021) does model a negative correlation between layoff risk and earnings, but that approach would not capture the fact that the layoff rate stabilizes after a certain earnings level. Robin (2011) allows for a form of differential layoff rates for high and low earners following an aggregate shock, but cannot produce such a pattern during "normal" times. Matching this pattern is unique to my model. Furthermore, the presence of this pattern holds important policy implications which I explore in the next section, as it implies some workers with low initial earnings can be stuck in an unstable zone of continued low earnings and more frequent churning into unemployment, while those who reach a higher earnings level are better shielded from future shocks. This speaks to the "unemployment scarring" literature, which observes that the best predictor of future unemployment is past unemployment (Arulampalam et al., 2001). A model without differential job destruction could not reproduce that pattern.

Figure 13: Quarterly transition rate into unemployment - model vs. data

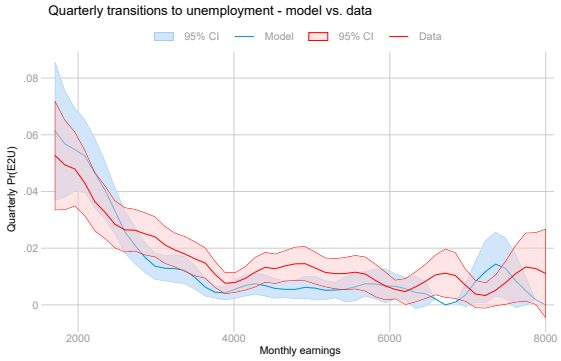

The fit of the model with the ratio of savings to earnings across the earnings distribution is plotted in Figure 14, with the distribution of those ratios in Figure 15. Of course in a model without savings there could be no such figure. My model produces a pattern of savings that increase with earnings in a way that does match that data. It also accords with the conclusion in Lise (2013) that desired saving increases with earnings. The evidence also supports the proposition in Carroll & Samwick (1998) that workers having a target level of savings given earnings. Moreover, I match the dispersion in those targets. Conditional on earnings, there is a range of target savings levels, which the model rationalizes as reflecting heterogeneous preferences over risk. As shown in the illustration in Figure 12, given those preferences we should expect the variance of savings to be increasing with income, which Figure 15 seems to suggest.

Figure 14: Average savings per earnings level - model vs. data

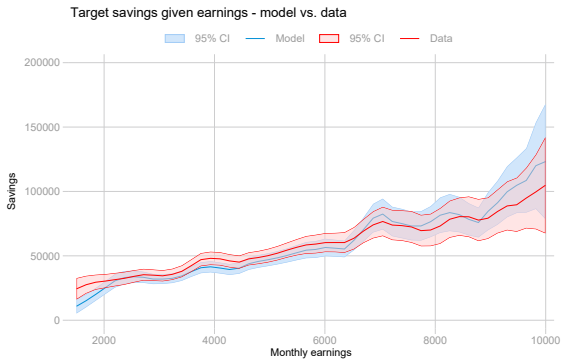

Figure 15: Savings per earnings ratios - model vs. data

Ideally, the distribution of risk preferences as I estimate it here will match other estimates from the literature. Chetty (2006) reports that the approach based on labor supply elasticities consistently finds an estimate of the mean coefficient of relative risk aversion around 1. Search model approaches that assume homogeneously risk averse workers (Lentz (2009); Rendon (2006); Lise (2013)) report results in the range of 1.2-2.2. Using bounded expected utility, Kimball et al. (2009) estimate a mean for the PSID sample of 3.8. I find a symmetrical distribution with a mean coefficient of relative risk aversion of 1.17, a median of 1.14, and a standard deviation of 0.49, which I plot in Figure 16. Inspecting the figure, it appears that if anything I slightly overestimate risk aversion because the workers simulated from the model save more on average given earnings than do the workers in the data. Thus my estimates generally match those of Chetty (2006) and are far from Kimball et al. (2009). One question to look into in the future is why these two approaches get results so far apart, and what the informational content is of the elicited risk preference survey questions in the PSID that yield such high values.

Figure 16: Estimated distribution of risk preferences

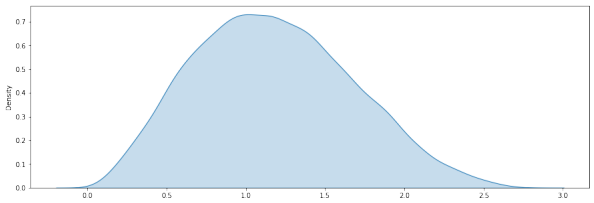

Because the model relies on several parametric assumptions and the use of order statistics, it is useful to know how sensitive the estimates are to these choices. I present the results from a series of these checks in Table 7. First, in the "direct" column I use a brute force optimization algorithm that searches the entire parameter space by dividing it into increasingly smaller hyperrectangles (Johnson, 2014). The new estimates are fairly close to the main ones, and all are within a standard deviation. Another issue to consider is whether the estimate of the flow benefit is influenced by sample trimming. One reassuring result is that my estimate of about 600 is close to the estimate obtained by Rendon (2006) of 650. To see how different estimates would influence the results, in the b +/- 100 columns I re-estimate the model by setting the flow benefit to 500 and 700. In general, the parameters related to shocks stay about the same, and the parameters related to risk preferences shift in a logical pattern. When the flow benefit is lower, the model rationalizes the same job acceptance rate and savings decisions by concluding that workers are less risk averse; when it is higher, a more risk averse distribution of preferences rationalizes the data. Third, in attempting to be flexible yet parametric, I model firm productivity and worker risk preferences as following beta distributions. For productivity this also leans on order statistics of the sample minimum and maximum earnings to obtain the scale for the draws. To avoid order statistics, the most manageable alternative to check is the lognormal distribution. In addition to being a common parametric assumption for wage offers, this fits with the finding in (Cunningham et al., 2019) of long right tails: the 99-90 firm productivity ratio is approximately equal to the 75-25 ratio. For risk preferences, Kimball et al. (2009) impose a lognormal distribution to obtain their estimates of risk tolerance. When I switch to lognormal distributions and rerun the model (column "lognormal"), I obtain very similar shock parameters. For risk preferences, this and the main specification find a mean of 1.17. Wage offers under the main specification have a mean of $2200 and a standard deviation of 708, while under the lognormal model they have a mean of $2552 and a standard deviation of 684. If anything, the more restrictive parametric assumptions fit the data slightly better based on the value of the objective function at the minimum. In contrast, reducing risk preferences to a single value for all workers results in a much poorer fit, as shown in column "1 risk". Several additional sensitivity checks are presented in Appendix B. Note that approaches that relax parametric assumptions for the wage offer distribution in particular, such as Bontemps et al. (2000), rely on modeling assumptions for identification. In particular, they assume that the set of accepted offers is the distribution of offers, and can therefore fit that observed distribution directly. When some offers are rejected, I need to assume some shape for the unobserved part of the distribution to get identification.

Table 7: Sensitivity of estimates to identifying assumptions

parameter	main	direct	b-100	b+100	lognorm	1 risk
α_x	3.283	5.895	1.496	5.049	—	4.824
β_x	9.315	13.522	5.855	12.212	—	12.249
λ_0	0.186	0.267	0.227	0.250	0.205	0.203
λ_1	0.086	0.074	0.113	0.088	0.061	0.086
θ	0.060	0.054	0.066	0.061	0.047	0.068
δ_0	0.160	0.147	0.160	0.141	0.164	0.212
η	0.050	0.044	0.061	0.063	0.028	0.062
α_γ	3.033	5.050	3.141	3.126	—	—
β_γ	4.714	10.376	5.346	4.178	—	—
b	604.974	604.974	500.000	700.000	604.974	604.974
μ_x	—	—	—	—	8.316	—
σ_x	—	—	—	—	0.267	—
μ_γ	—	—	—	—	0.078	—
σ_γ	—	—	—	—	0.696	—
γ	—	—	—	—	—	0.728
fval	1929.767	2144.083	1892.882	2127.347	1335.956	4016.951

In this section I combined the model with the PSID data to estimate the structural parameters governing labor market dynamics and worker behavior. I analyzed the results, checking how they fit the data, how they compared to other estimates in the literature, and how sensitive they were to some of the identifying assumptions I made. In the next section I will use these parameters to conduct a counterfactual policy analysis and to simulate the impact of an aggregate downturn.

7 Policy experiments and downturn simulation

In this section I use the model to conduct counterfactual analysis. First, I use the parameters I estimated in the previous section to compare the cost, welfare, and labor market effect of three social insurance policies. Then I simulate a downturn as described in Section 3 and consider how the predictions of the model depend on accumulated precautionary savings. Finally, I combine the two analyses by studying the impact of the three social insurance policies during a downturn.

Typically, economists look for a market imperfection or inefficient outcome as motivation for policy intervention. For instance, Acemoglu & Shimer (2000) argue that while unemployment insurance reduces employment, it shifts workers into more productive jobs that they would not otherwise desire, potentially raising aggregate output. In the environment here, policies that raise the reservation wage should likewise reallocate workers into more productive jobs, as productivity is positively linked to pay, with the difference being that workers already prefer those positions. A stronger policy motivation is equity. Even if the prevailing dynamics prove optimal in aggregate, there are clear winners and losers among otherwise identical agents, and no inherent mechanism by which the winners compensate the losers. The position assumed here is that a government has opted for intervention, and the question is which type of intervention produces the largest welfare gains relative to program costs.

The three policies I study are 1) unemployment insurance, 2) a one-time cash transfer, and 3) an earnings subsidy that also covers zero earnings. To consider extended unemployment benefits, I also re-run the unemployment insurance experiment with an additional benefit of $300 per month. My model is uniquely well positioned to run these experiments because of two features: first, because it has savings I can look at a cash transfer without assuming people have to spend the entire amount immediately, and because I allow for different preferences I can allow for different degrees of responsiveness to transfer. Second is the related

fact that I can actually look at consumption, which is a closer proxy for welfare than income alone. The uniqueness of the PSID data let me estimate a model with consumption in the previous section, and having consumption separate from in the model is invaluable for actually getting at welfare. This will be doubly useful later on in the analysis of the role of savings in a downturn, as shocks that affect the labor market often affect the stock of wealth as well.

Unemployment insurance is a common policy to study with a search model. Such a policy should raise consumption for those out of work and enable them to be more patient and wait for better job offers. The other side of that coin, of prominence in policy debates during the Covid-19 pandemic, is that this greater patience means that the unemployment rate will be higher than it would be otherwise, which many policymakers find an undesirable outcome. While one way to reduce unemployment in that scenario is to reduce the unemployment benefit, this may run counter to the original policy objective. The U.S. government also provided large stimulus checks to most Americans to support them during the pandemic, but outside of Rendon (2006) I am not aware of any papers that evaluate a large payment like that in the context of a dynamic model of the labor market, and that paper only looks at an initial transfer upon first starting out in the labor market. The third alternative is an earnings subsidy that constitutes a gradual tapering off of unemployment benefits for people who accept new jobs instead of taxing benefits at 100% when someone resumes work. This type of policy is also uncommon in counterfactual analysis though it has received more attention in other strands of the literature (see Phelps (1997) for an argument in favor of wage subsidies, and Moffitt (2002) for a discussion and summary of the evidence, with a focus on the earned income tax credit in the United States). A concern with subsidizing a range of earnings is that workers may drop down from higher paying jobs, or stay put and reject higher offers. Here I am not attempting to find the optimal subsidy level and taper rate to balance a government budget or achieve a certain welfare outcome. Rather, the goal is to demonstrate the relative consumption and employment impact, and relative costs, of the policy options at a reasonable implementation level. For instance, the baseline unemployment insurance rate is approximately the monthly average in the United States, and the subsidy is chosen to bring about the equivalent of a $15/hour full-time minimum wage.

I present the results from simulating each of the three policies in Table 8 up to two years after implementation in the steady state. The specific values I use are $1000 per month for unemployment insurance, $5000 for the cash transfer, and a taper rate of $0.75 per dollar of earnings above $1300 for the earnings subsidy. That is, earnings below $1300, including $0, receive a $1000, and the subsidy is reduced by $0.75 for each dollar of earnings after that. As expected, unemployment benefits raise the unemployment rate relative to the baseline. Even though it also pays the same unemployment benefit, the low-income earnings subsidy has a much smaller impact on the unemployment rate, which by the eighth quarter rises by one percentage point instead of over four, as the marginal tax rate on benefits after accepting a job is not automatically 100%. Workers can accept lower paying jobs and search on the job rather than wait in unemployment for a high offer. Consumption at the 10th quantile is 14% higher than in the baseline scenario, and the 90-10 ratio of consumption, a proxy for inequality, falls by 9%. As a rough proxy for poverty, I calculate the share of workers earning less than $2000 per month, which is between the U.S. poverty thresholds for a household with one or two children. Unemployment insurance does not affect this much. The earnings subsidy cuts it in half, from 5.9% to 2.9% after eight quarters. As for the one-time transfer, that has a short run increase in consumption and a decrease in the poverty proxy, yet little impact on unemployment and no discernible long run effect. To me this suggests that even though workers could choose to save the transfer, many of them are already at or near their target level of savings and so consume most of it quickly to regain that

target. Consumption at the 10th quantile is about 6% higher in the first few quarters after the transfer, and by two years out it remains just 3% above the baseline level.

Table 8: Effect of social insurance policies

Outcome	Policy	Quarter 1	2	3	4	6	8
Consumption	Baseline	100.00	100.00	100.00	100.00	100.00	100.00
10th percentile	UI	108.71	107.38	106.92	106.57	106.72	106.50
	UI+300	110.28	108.51	107.68	107.26	106.91	105.97
	Subsidy	114.19	112.97	112.16	111.42	111.16	110.89
	Transfer	105.00	104.57	103.95	103.68	103.21	102.66
Consumption	Baseline	0.049	0.049	0.051	0.049	0.052	0.054
% under $2000	UI	0.037	0.041	0.046	0.049	0.053	0.057
	UI+300	0.029	0.031	0.034	0.037	0.045	0.055
	Subsidy	0.018	0.020	0.024	0.024	0.026	0.027
	Transfer	0.033	0.036	0.040	0.041	0.046	0.049
Unemployment	Baseline	0.033	0.032	0.031	0.030	0.029	0.030
	UI	0.041	0.046	0.051	0.053	0.059	0.063
	UI+300	0.043	0.050	0.055	0.060	0.068	0.075
	Subsidy	0.035	0.035	0.036	0.035	0.035	0.036
	Transfer	0.033	0.032	0.032	0.031	0.030	0.030
Inequality	Baseline	1.000	1.000	1.000	1.000	1.000	1.000
90/10 ratio	UI	0.950	0.960	0.962	0.965	0.961	0.961
	UI+300	0.945	0.958	0.963	0.967	0.967	0.973
	Subsidy	0.924	0.932	0.938	0.943	0.942	0.941
	Transfer	0.964	0.967	0.972	0.974	0.978	0.981

Note: 10th percentile of consumption and inequality indexed to baseline.

Not only does the added earnings subsidy affect unemployment less than pure unemployment insurance, but as shown in Table 9 the relative rise in employment is eventually enough such that after two years, total monthly program expenditures are nearly the same even though the policy reaches twice as many workers. Initially, the subsidy costs nearly twice as much because the unemployment rate has not had the chance to rise to its new steady state level, yet by two years out the subsidy costs just 14% more, and the gap is shrinking as unemployment adjusts to the new benefit level. Relative to a $300 extended unemployment benefit, the subsidy ends up costing 4% less by the eighth quarter, with a 4 point lower unemployment rate. The transfer is the most expensive policy by far, exceeding the cost of unemployment insurance by a multiple of 1200. In this case, I allocated the transfer to every worker. It would be cheaper if I imposed an income cap, but still pricier than the other policies relative to its effect on employment.

Table 9: Social insurance policy costs and population coverage

	Policy	Quarter 1	2	3	4	6	8
Cost	UI	100.00	100.00	100.00	100.00	100.00	100.00
	UI+300	103.95	107.33	108.28	111.80	116.19	119.48
	Subsidy	176.34	154.73	140.84	131.63	121.32	114.55
	Transfer	12082.72	0.00	0.00	0.00	0.00	0.00
Coverage	UI	0.04	0.05	0.05	0.05	0.06	0.06
	UI+300	0.04	0.05	0.06	0.06	0.07	0.08
	Subsidy	0.16	0.16	0.16	0.16	0.16	0.16
	Transfer	1.00	0.00	0.00	0.00	0.00	0.00

Costs indexed to cost of baseline UI benefit.

Next, I repeat the exercise while simultaneously simulating a downturn.[24] First, I discuss its anticipated

[24]Technically, in the model this is a permanent negative shock, not a cyclical recession. I abuse the term here for illustrative

impact and calibrate an appropriate magnitude. A downturn in the model reduces profit margins for prevailing jobs, which should send a greater fraction of employed workers into unemployment each period. These workers are more likely to be from lower productivity, and therefore lower wage, matches. Accordingly, the unemployment rate should rise, and there should be more frequent churning into unemployment. Since wage offers for new matches are be lower, this should lead to slower job productivity growth for the workers beginning to climb back up the ladder. Wage inequality may widen, as the initial highest earners (with the corresponding lowest risk) are most likely to keep their current jobs. Furthermore, re-employed "displaced" workers could earn lower wages even at similar looking firms than identical, consistently employed counterparts.

How does these implications align with what occurred during recent actual downturns? Examining income losses by displaced workers in the Great Recession, Lachowska et al. (2020) find that "displaced workers' earnings losses occurred mainly because hourly wage rates dropped at the time of displacement." When comparing them to non-displaced workers, they observe, "five years after job loss, displaced workers' earnings were 16 log points less than those of a stably employed comparison group." As additional evidence of the unequal impact, Figure 17 plots wage and consumption inequality over the two-year intervals available for the PSID sample. The 90-10 wage ratio increases by about 25% between 2008 and 2010, and the 90-10 consumption ratio rises about 15% over its 2006 low point. Both fall as the recovery begins yet remain above their pre-recession levels.

Figure 17: A rise in wage and consumption inequality during the Great Recession

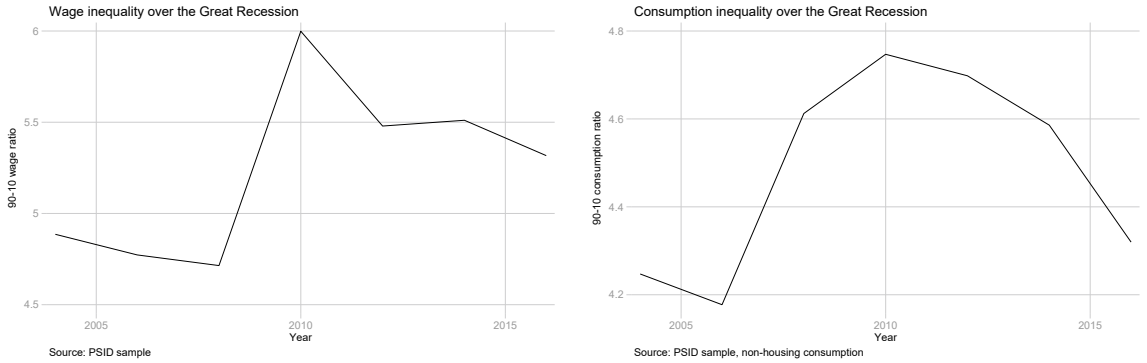

Some additional facts from the Great Recession that can help calibrate the reasonableness of the modeled downturn are that the unemployment rate approximately doubled, rising from 5% to 10% in 2008-2009, and as shown in Figure 18 mean non-housing consumption dropped by slightly over 8% for the PSID sample. As for the magnitude of the downturn, one useful metric is the "output gap", which is the difference between actual and potential output as estimated by the Congressional Budget Office. The output gap in the United States reached 5.5% in 2009 and 10.5% in 2020.[25] To calibrate the magnitude to the output gap observed during the Great Recession, I set s to shock every match product downwards by 500 units–an amount equivalent to about 6% of mean productivity.

purposes.

[25] `https://fred.stlouisfed.org/graph/?g=Ffx7`

Figure 18: Mean consumption over the Great Recession

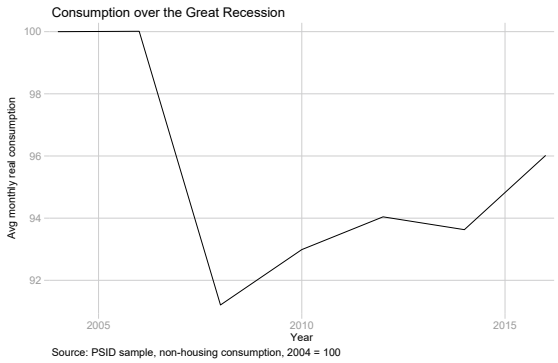

For simplicity of exposition and to focus on the dynamics of a downturn, I model the downturn here as affecting the entire economy, though one could readily imagine this as the scenario affecting a specific local labor market or sector. As an example, consider the local labor market impact of the China Shock of the early 2000s. The reduced form empirical results from the China Shock literature (Autor et al., 2016) describe the local labor market impact as follows: "Adjustment in local labor markets is remarkably slow, with wages and labor-force participation rates remaining depressed and unemployment rates remaining elevated for at least a full decade after the China trade shock commences. Exposed workers experience greater job churning and reduced lifetime income."

Figure 19 shows that the simulated downturn matches the unequal expansion in layoff risk for the lowest earners as seen in the Great Recession and Covid-19 pandemic depicted in Figure 2.

Figure 19: Quarterly transition rate into unemployment before and after downturn

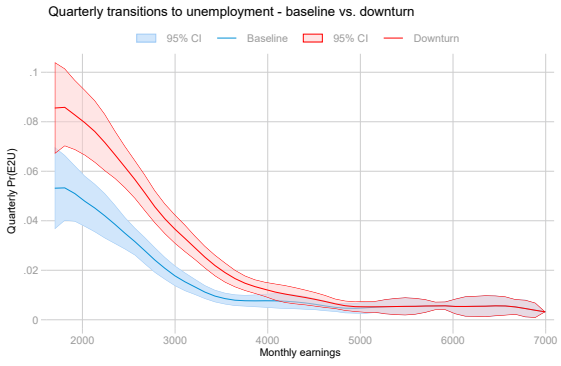

In the simulation, the new prevailing distribution of earnings offers shifts downward. I compare the scenarios in Figure 20, where it is clear that workers who accept jobs during the downturn are accepting much worse offers than before. In other words, one reason workers may remain at their current job even when the layoff risk rises is that their outside option has also worsened.

Figure 20: Accepted earnings offer distribution before and after downturn

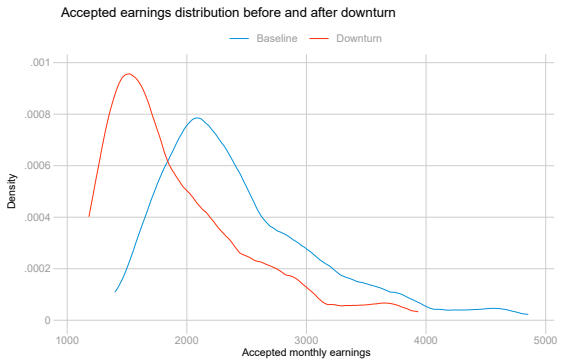

To see how the predictions of the model vary with the level of savings across the distribution of consumption, first consider the baseline impact of erasing initial savings, depicted in Table 10. The table shows how each percentile of the consumption distribution changes over up to eight quarters when I wipe out initial savings in the model, relative to a scenario where initial savings are unchanged. As would be expected since savings are lowest for the least well off workers, the relative decline in the 10th percentile of consumption is lower than for the other percentiles, dropping by 9% initially compared to a drop of about 12-13% for higher consumers. Each percentile rises over time as workers build back up their buffer stocks.

Table 10: Impact of erasing initial savings on the distribution of consumption - normal scenario

	consumption percentile				
quarter	10th	25th	50th	75th	90th
0	100.00	100.00	100.00	100.00	100.00
1	90.87	86.99	88.04	88.49	86.93
2	91.51	88.00	89.05	89.35	88.13
3	92.30	88.92	89.90	90.02	89.09
4	92.55	89.89	90.69	90.66	89.82
6	93.15	91.18	91.89	91.54	90.86
8	93.51	92.20	92.69	92.38	91.76

When we compare the consumption response after eliminating initial savings to what the model otherwise predicts for a downturn, the implications change. Table 11 shows the cross-section of consumption levels during a downturn with zero initial savings relative to the prevailing levels under a downturn with initial savings unchanged. In the first quarter, we see the same pattern as in the normal scenario in Table 10, whereby the 90th percentile of consumption drops the most and the 10th percentile the least during a downturn that also eliminates accumulated wealth. However, by the 4th quarter of the downturn the pattern is reversed: the largest relative drop is for the 10th percentile of consumption. This is because in the simulated downturn, as shown below in Table 13, the unemployment rate rises. In addition, two model features become important: first, new wages offers are lower, leading to lower accepted wages (and thus less income available for consumption); second, layoff risk for lower earners is particularly higher, leading to higher desired precautionary savings (and thus lower consumption when below the target level). When people do not have savings, they take worse jobs, and in addition to paying lower wages those jobs are the most unstable.

Table 11: Impact of erasing initial savings on the distribution of consumption - downturn scenario

	consumption percentile				
quarter	10th	25th	50th	75th	90th
0	100	100	100	100	100
1	91.54	89.28	89.00	88.95	87.52
2	89.94	90.10	89.84	89.91	88.53
3	87.76	90.84	90.49	90.55	89.32
4	85.38	91.80	91.01	91.16	89.96
6	86.49	92.52	91.90	92.11	91.16
8	88.79	92.66	92.46	92.78	92.03

Next, in Table 12 I expand to showing how the downturn on its own and in conjunction with a loss of savings affects a broader set of outcomes related to worker welfare. In the simulation, as in the Great Recession, the unemployment rate roughly doubles within the first two quarters and nearly triples after two years. We can also see that my proxy for the poverty rate rises by 50% within a few quarters and doubles after two years. By then, consumption inequality has risen about 60% compared to the steady state counterfactual, and the 10th percentile of consumption is down one-third while overall mean consumption falls by 5%. Here we also see the importance of savings. In the no savings scenario, I suppose that the downturn also wipes out accumulated wealth. When that happens, mean consumption drops by about $600, or 8% more than otherwise in the downturn, and 13% relative to the normal scenario. Consumption inequality is about 4% higher. The poverty proxy rises by several percentage points, which translates to a nearly 33% higher level immediately. The unemployment rate is actually slightly lower when the downturn wipes out savings, because workers without savings have a lower reservation wage and are accepting worse offers. The downturn reduces mean wage accepted out of unemployment by about 18%, and the 10th quantile of accepted wage offers by roughly the same amount. When the downturn also wipes out savings, wages accepted by unemployed workers fall even further–closer to 28%. for the 10th percentile in Quarter 1, and 23% at the mean.

Table 12: Simulated impact of an aggregate downturn

		Quarter					
Outcome	Policy	1	2	3	4	6	8
Consumption	Baseline	100.00	100.00	100.00	100.00	100.00	100.00
Mean	Downturn	96.62	96.15	95.80	95.49	95.02	94.71
	No initial savings	84.86	85.25	85.63	85.96	86.50	87.05
Consumption	Baseline	100.00	100.00	100.00	100.00	100.00	100.00
10th percentile	Downturn	95.66	90.56	84.29	77.38	70.07	67.62
	No initial savings	87.11	80.73	72.02	66.07	61.45	59.77
Consumption	Baseline	0.049	0.049	0.051	0.049	0.052	0.054
% under $2000	Downturn	0.069	0.085	0.098	0.110	0.129	0.144
	No initial savings	0.088	0.105	0.120	0.131	0.152	0.167
Inequality	Baseline	1.00	1.00	1.00	1.00	1.00	1.00
90/10 ratio	Downturn	1.03	1.09	1.17	1.28	1.41	1.47
	No initial savings	0.98	1.07	1.21	1.33	1.45	1.51
Accepted wage	Baseline	100.00	100.00	100.00	100.00	100.00	100.00
Mean	Downturn	78.39	79.71	77.64	77.74	80.93	80.22
	No initial savings	74.12	75.91	74.83	74.77	77.76	77.13
Accepted wages	Baseline	100.00	100.00	100.00	100.00	100.00	100.00
10th percentile	Downturn	79.08	82.04	76.96	77.38	78.45	76.93
	No initial savings	70.13	73.15	71.44	70.27	74.55	75.20
Unemployment	Baseline	0.033	0.032	0.031	0.030	0.029	0.030
	Downturn	0.046	0.054	0.059	0.060	0.067	0.070
	No initial savings	0.044	0.051	0.055	0.055	0.062	0.065

Note: consumption, inequality, and wages indexed to baseline.

Next, I implement the same three counterfactual policies starting at the onset of the simulated downturn. The results are in Table 13. Here, the relative labor market effect of the earnings subsidy is more pronounced. Under the subsidy, the unemployment rate is actually slightly lower than without any policy intervention. By contrast, with unemployment insurance alone, the gap rises over time, reaching three percentage points after a year and five points after two years. The subsidy cuts the post-transfer proxy poverty rate significantly, from 6.8% to 2.2% in the first quarter of the downturn. Consumption for the 10th quantile is 31% higher under the earnings subsidy a year into the downturn, compared to 10% under unemployment insurance. The 90-10 consumption ratio is 19% lower after a year than if there were no intervention. Meanwhile, the effect of the asset transfer is again short-lived. It props up consumption for the first quarter or two, and then the impact dissipates and the scenario begins to resemble the one with no intervention. If the policy goal is to help people get better jobs, unemployment insurance is more effective, as it leads to much higher accepted wages at the 10th percentile—22% higher than with no intervention. This is because people can afford to wait in unemployment until they receive a good offer. The subsidy lowers accepted wages slightly, as low paying jobs become attractive when combined with the subsidy income. A subsidy does not help workers find better jobs; it makes bad jobs pay better.

Table 13: Effect of social insurance policies in a downturn

		Quarter					
Outcome	Policy	1	2	3	4	6	8
Consumption	Downturn	100.00	100.00	100.00	100.00	100.00	100.00
10th percentile	Downturn UI	109.39	108.90	111.36	116.20	118.71	112.78
	Downturn UI+300	111.55	111.88	115.62	121.49	126.48	126.14
	Downturn subsidy	115.23	119.58	126.80	135.79	147.43	151.17
	Downturn transfer	105.09	105.59	106.26	105.72	104.20	103.60
Consumption	Downturn	0.069	0.085	0.098	0.110	0.129	0.144
% under $2000	Downturn UI	0.050	0.063	0.076	0.085	0.101	0.114
	Downturn UI+300	0.037	0.044	0.052	0.060	0.075	0.092
	Downturn subsidy	0.024	0.030	0.035	0.037	0.042	0.044
	Downturn transfer	0.055	0.073	0.090	0.102	0.123	0.138
Unemployment	Downturn	0.046	0.054	0.059	0.060	0.067	0.070
	Downturn UI	0.054	0.069	0.084	0.095	0.116	0.132
	Downturn UI+300	0.055	0.072	0.087	0.101	0.124	0.142
	Downturn subsidy	0.045	0.051	0.054	0.054	0.059	0.061
	Downturn transfer	0.046	0.054	0.060	0.061	0.068	0.071
Inequality	Downturn	1.000	1.000	1.000	1.000	1.000	1.000
90/10 ratio	Downturn UI	0.947	0.950	0.927	0.887	0.864	0.906
	Downturn UI+300	0.941	0.937	0.904	0.858	0.821	0.819
	Downturn subsidy	0.931	0.894	0.841	0.783	0.716	0.694
	Downturn transfer	0.965	0.959	0.951	0.956	0.968	0.973
Accepted wages	Downturn	100.000	100.000	100.000	100.000	100.000	100.000
10th percentile	Downturn UI	101.656	103.781	106.355	110.255	117.514	124.507
	Downturn UI+300	101.859	104.376	107.258	111.489	119.528	127.504
	Downturn subsidy	99.561	99.430	98.606	97.609	96.405	94.031
	Downturn transfer	100.063	100.155	100.162	100.377	100.294	100.626

Note: consumption, wages, and inequality indexed to baseline.

Relative costs of the social insurance programs are listed in Table 14. Compared to the baseline $1000 unemployment benefit, the expanded $1300 benefit costs 5% more after four quarters and 8% more after eight. The combined earnings subsidy and unemployment benefit costs 30% more after eight quarters and 16% more after eight. Again, the transfer is an outlier, costing several orders of magnitude more than the monthly cost of the other policies.

Table 14: Social insurance policy costs and population coverage in a downturn

		Quarter					
	Policy	1	2	3	4	6	8
Cost	Downturn UI	100.00	100.00	100.00	100.00	100.00	100.00
	Downturn UI+300	101.57	103.50	104.11	105.67	106.62	107.92
	Downturn subsidy	167.12	149.17	137.07	130.09	120.51	115.82
	Downturn transfer	9282.90	0.00	0.00	0.00	0.00	0.00
Coverage	Downturn UI	0.05	0.07	0.08	0.10	0.12	0.13
	Downturn UI+300	0.06	0.07	0.09	0.10	0.12	0.14
	Downturn subsidy	0.17	0.18	0.19	0.20	0.21	0.22
	Downturn transfer	1.00	0.00	0.00	0.00	0.00	0.00

Costs indexed to cost of baseline UI benefit.

8 Conclusion

The objective of the analysis has been to examine how social insurance policies affect the welfare and employment of low-income workers during an economic downturn. In particular, it considers the role of additional unemployment benefits and savings in the job offer acceptance decision. Comparing several policy options intended to raise welfare for low-income workers, I find that one-off transfers have a transitory consumption effect without much altering employment. Raising unemployment benefits by $300 raises unemployment by less than a percentage point. An alternative policy of combining unemployment insurance with a tapered earnings subsidy raises consumption and reduces inequality without the employment deterring incentives inherent in unemployment insurance alone. Inequality under the joint policy is 24% lower two years into a downturn, and 30% lower than with no intervention. In a normal economy employment is 3.3 percentage points higher under the joint policy than under unemployment insurance alone. One year into a downturn the combined policy yields a 3.3 percentage point lower unemployment rate, and after two years the employment gain is 7 percentage points. Furthermore, the corresponding higher employment rate may be such that the the total additional costs of the larger policy are small. Two years after implementation, the difference is only 14% during normal times and 16% during a downturn, and only 7% more relative to a policy of expanded unemployment benefits. Both unemployment insurance and an earnings subsidy are less costly than large-scale one-off transfers, which as calibrated equal 600-1200 times the monthly price of the flow benefits.

The model developed to support this analysis offers a clear mechanism for why job destruction rates are observed to differ systematically with earnings during the steady state and to do so more sharply in an economic downturn. It implies that low earnings are self-reinforcing because they also increase the risk of returning to unemployment and starting over on the job ladder. This dynamic tends to increase inequality, especially during an economic downturn. The mechanism ties together related observations from the literature on firm wage premiums, productivity, and match duration, and also contributes to the literature suggesting that given sufficient search frictions, compensating differentials may not be expected to arise. For job security this is interesting because any demand for compensation for bearing additional layoff risk implicitly relies on search frictions being present. Absent such frictions, there is no consumption loss associated with job termination, as an identical position can be obtained immediately.

One implication of the relationship between earnings and layoff risk characterized here is that even when workers have different preferences over risk, those differences will not induce sorting behavior across jobs. The result obtains in the model because the same types of matches are attractive to workers with a relatively high utility weighting of consumption and to those with a relatively strong preference for stability.

An assumption underpinning this result is that all matches face the same distribution of shocks. If the model were expanded to two sectors differing in the frequency or distribution of productivity shocks—and hence the chance of layoffs—risk preferences might prove more salient for worker labor market behavior. However, empirical approaches to testing for sorting on the basis of risk preferences tend to rely on survey responses such as those in the PSID, which I have observed here to not align well with preferences as revealed through savings behavior. Given the contradiction, a first next step is to better understand the alignment between stated and revealed risk preferences.

A second implication of the mechanism I propose is that *given* match productivity, there *is* an inherent tradeoff between layoff risk and earnings. Following a negative shock, a firm that is paying different wages to two identical workers has a profit motive to lay off the higher paid worker first. The observation yields a direct empirical prediction when combined with the literature on gender differences in bargaining power over match surplus (Card et al., 2016). If women tend to receive a smaller share of the surplus than men in equivalent matches, then all else being equal women should experience lower rates of displacement. Although the reasoning is not the same as that suggested by Sahin et al. (2010), it is consistent with their observation that male and female unemployment rates diverged from 5.1% and 4.9% respectively in 2007 to 11% and 8.3% during the recession in 2009.

From a policy perspective, there is an inherent tension between the goals of supporting unemployed workers so that they can afford to wait for better job offers and of reducing the unemployment rate, particularly in a bad economic state. The counterfactual experiments in the paper suggest a less considered alternative to the comparison between providing and not providing unemployment insurance. On the surface, adding an earnings subsidy on top of unemployment insurance is unambiguously a more comprehensive policy intervention. However, because it provides strong incentives for employment, a tapered subsidy may reduce the unemployment rate relative to a policy of unemployment insurance alone by a margin large enough that the total costs of each intervention become comparable in magnitude.[26]

In building a model equipped to conduct this policy analysis, I make a number of simplifying assumptions which I endeavor to show are not essential to the results obtained. My conclusion from these exercises is that the most important assumption is that wages and profitability from a match rise in tandem. In that way, higher paying jobs are associated with a higher buffer such that the match can absorb larger shocks before it turns unprofitable. This assumption has empirical support. For instance, Autor et al. (2020) observe that the most productive firms within an industry have high markups and a low labor share of value added. Despite a lower labor share they detect higher average wages, such that firms and workers both benefit from a larger pie.

A broader question involves the concept of a more general equilibrium in which firms and workers can freely adjust their behavior in response to the actions of each other and any policy interventions. While this is an important feature conceptually, as it is not central to the job destruction dynamics that are the focus here, I follow the literature closest to the current paper in abstracting away from such considerations. With that caveat noted, the results here suggest potentially important equity consequences deriving from the relationship between layoff risk and earnings. Rather than increasing unemployment benefits or issuing large checks, a policymaker concerned with raising employment and reducing consumption inequality might find a cost effective tool in a low-earnings subsidy.

[26] I abstract away from implementation issues. Evaluation of the UK's Universal Credit finds harmful effects from lengthy delays in the registration process and from stringent participation requirements (Craig & Katikireddi, 2020).

References

Abowd, J. M., Kramarz, F., & Margolis, D. N. (1999). High wage workers and high wage firms. Econometrica, 67(2), 251–333.

Acemoglu, D., & Shimer, R. (1999). Efficient unemployment insurance. Journal of Political Economy, 107(5), 893–928.

Acemoglu, D., & Shimer, R. (2000). Productivity gains from unemployment insurance. European economic review, 44(7), 1195–1224.

Altonji, J. G., & Devereux, P. J. (1999). The extent and consequences of downward nominal wage rigidity. Tech. rep., National Bureau of Economic Research.

Arulampalam, W., Gregg, P., & Gregory, M. (2001). Introduction: unemployment scarring. The Economic Journal, 111(475), F577–F584.

Autor, D. (2021). The labor shortage has empowered workers. The New York Times.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. The Quarterly Journal of Economics, 135(2), 645–709.

Autor, D. H., Dorn, D., & Hanson, G. H. (2016). The china shock: Learning from labor-market adjustment to large changes in trade. Annual Review of Economics, 8, 205–240.

Aw, B. Y., Chen, X., & Roberts, M. J. (2001). Firm-level evidence on productivity differentials and turnover in taiwanese manufacturing. Journal of Development Economics, 66(1), 51–86.

Azariadis, C. (1975). Implicit contracts and underemployment equilibria. Journal of Political Economy, 83(6), 1183–1202.

Baily, M. N., Bartelsman, E. J., & Haltiwanger, J. (2001). Labor productivity: structural change and cyclical dynamics. Review of Economics and Statistics, 83(3), 420–433.

Barattieri, A., Basu, S., & Gottschalk, P. (2014). Some evidence on the importance of sticky wages. American Economic Journal: Macroeconomics, 6(1), 70–101.

Bewley, T. F. (2009). Why wages don't fall during a recession. Harvard university press.

Bhutta, N., Dettling, L. J., et al. (2018). Money in the bank? assessing families' liquid savings using the survey of consumer finances. Tech. rep., Board of Governors of the Federal Reserve System (US).

Bonhomme, S., & Jolivet, G. (2009). The pervasive absence of compensating differentials. Journal of Applied Econometrics, 24(5), 763–795.

Bontemps, C., Robin, J.-M., & Van den Berg, G. J. (2000). Equilibrium search with continuous productivity dispersion: Theory and nonparametric estimation. International Economic Review, 41(2), 305–358.

Burdett, K., Carrillo-Tudela, C., & Coles, M. (2020). The cost of job loss. The Review of Economic Studies, 87(4), 1757–1798.

Burdett, K., & Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. International Economic Review, (pp. 257–273).

Burgess, S., & Turon, H. (2010). Worker flows, job flows and unemployment in a matching model. European Economic Review, 54(3), 393–408.

Cahuc, P., Postel-Vinay, F., & Robin, J.-M. (2006). Wage bargaining with on-the-job search: a structural econometric model. Econometrica, 74(2), 323–364.

Caldwell, S., & Harmon, N. (2019). Outside options, bargaining, and wages: Evidence from coworker networks. Unpublished manuscript, Univ. Copenhagen.

Campbell III, C. M., & Kamlani, K. S. (1997). The reasons for wage rigidity: evidence from a survey of firms. The Quarterly Journal of Economics, 112(3), 759–789.

Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. Journal of Labor Economics, 36(S1), S13–S70.

Card, D., Cardoso, A. R., & Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. The Quarterly Journal of Economics, 131(2), 633–686.

Carroll, C. D., Dynan, K. E., & Krane, S. D. (2003). Unemployment risk and precautionary wealth: Evidence from households' balance sheets. Review of Economics and Statistics, 85(3), 586–604.

Carroll, C. D., & Samwick, A. A. (1998). How important is precautionary saving? Review of Economics and Statistics, 80(3), 410–419.

Chetty, R. (2006). A new method of estimating risk aversion. American Economic Review, 96(5), 1821–1834.

Couch, K. A., & Placzek, D. W. (2010). Earnings losses of displaced workers revisited. American Economic Review, 100(1), 572–89.

Craig, P., & Katikireddi, S. V. (2020). Early impacts of universal credit: the tip of the iceberg? The Lancet Public Health, 5(3), e131–e132.

Cunningham, C., Foster, L., Grim, C., Haltiwanger, J., Pabilonia, S. W., & Stewart, J. (2019). Dispersion in dispersion: Measuring establishment-level differences in productivity. Center for Economic Studies, CES 18-25R.

Dalton, M. R., & Groen, J. A. (2020). How do jobseekers search for jobs? new data on applications, interviews, and job offers. Beyond the Numbers: Employment & Unemployment, 9, no. 14.

Danforth, J. P. (1979). On the role of consumption and decreasing absolute risk aversion in the theory of job search. Studies in the Economics of Search, 123, 109–131.

Dey, M. S., & Flinn, C. J. (2005). An equilibrium model of health insurance provision and wage determination. Econometrica, 73(2), 571–627.

Dorn, D., Katz, L. F., Patterson, C., Van Reenen, J., et al. (2017). Concentrating on the fall of the labor share. American Economic Review, 107(5), 180–85.

Elsby, M. W., & Solon, G. (2019). How prevalent is downward rigidity in nominal wages? international evidence from payroll records and pay slips. Journal of Economic Perspectives, 33(3), 185–201.

Fackler, D., Mueller, S., & Stegmaier, J. (2021). Explaining wage losses after job displacement: Employer size and lost firm wage premiums. Journal of the European Economic Association.

Fallick, B. C., Lettau, M., & Wascher, W. L. (2020). Downward nominal wage rigidity in the united states during and after the great recession.

Farber, H. S. (2015). Job loss in the great recession and its aftermath: Us evidence from the displaced workers survey. Tech. rep., National Bureau of Economic Research.

Farber, H. S. (2017). Employment, hours, and earnings consequences of job loss: Us evidence from the displaced workers survey. Journal of Labor Economics, 35(S1), S235–S272.

Fehr, E., & Falk, A. (1999). Wage rigidity in a competitive incomplete contract market. Journal of Political Economy, 107(1), 106–134.

Feinberg, R. M. (1977). Risk aversion, risk, and the duration of unemployment. The Review of Economics and Statistics, (pp. 264–271).

Flinn, C., & Heckman, J. (1982). New methods for analyzing structural models of labor force dynamics. Journal of Econometrics, 18(1), 115–168.

Flood, S., King, M., Rodgers, R., Ruggles, S., & Warren, J. R. (2018). Integrated public use microdata series, current population survey: Version 6.0.

Foster, L., Haltiwanger, J., & Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? American Economic Review, 98(1), 394–425.

Francesconi, M., & Van der Klaauw, W. (2007). The socioeconomic consequences of "in-work" benefit reform for british lone mothers. Journal of Human Resources, 42(1), 1–31.

Friedrich, B., Laun, L., Meghir, C., & Pistaferri, L. (2019). Earnings dynamics and firm-level shocks. Tech. rep., National Bureau of Economic Research.

Gibbons, R., & Katz, L. F. (1991). Layoffs and lemons. Journal of Labor Economics, 9(4), 351–380.

Grigsby, J., Hurst, E., & Yildirmaz, A. (2021). Aggregate nominal wage adjustments: New evidence from administrative payroll data. American Economic Review, 111(2), 428–71.

Griliches, Z., & Regev, H. (1995). Firm productivity in israeli industry 1979–1988. Journal of Econometrics, 65(1), 175–203.

Grossman, G. M., & Oberfield, E. (2021). The elusive explanation for the declining labor share. Tech. rep., National Bureau of Economic Research.

Guiso, L., Pistaferri, L., & Schivardi, F. (2005). Insurance within the firm. Journal of Political Economy, 113(5), 1054–1087.

Haefke, C., Sonntag, M., & Van Rens, T. (2013). Wage rigidity and job creation. Journal of Monetary Economics, 60(8), 887–899.

Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. American Economic Review, 95(1), 50–65.

Hall, R. E., & Krueger, A. B. (2012). Evidence on the incidence of wage posting, wage bargaining, and on-the-job search. American Economic Journal: Macroeconomics, 4(4), 56–67.

Haltiwanger, J., Hyatt, H., & McEntarfer, E. (2018). Who moves up the job ladder? Journal of Labor Economics, 36(S1), S301–S336.

Haltiwanger, J. C., Hyatt, H. R., McEntarfer, E., & Staiger, M. (2021). Cyclical worker flows: Cleansing vs. sullying. Tech. rep., National Bureau of Economic Research.

Heathcote, J., Perri, F., & Violante, G. L. (2020). The rise of us earnings inequality: Does the cycle drive the trend? Review of Economic Dynamics, 37, S181–S204.

Herkenhoff, K. F. (2019). The impact of consumer credit access on unemployment. The Review of Economic Studies, 86(6), 2605–2642.

Hornstein, A., Krusell, P., & Violante, G. L. (2011). Frictional wage dispersion in search models: A quantitative assessment. American Economic Review, 101(7), 2873–98.

Hotz, V. J., & Scholz, J. K. (2007). 3. The Earned Income Tax Credit. University of Chicago press.

Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. The Quarterly Journal of Economics, 124(4), 1403–1448.

Hwang, H.-s., Mortensen, D. T., & Reed, W. R. (1998). Hedonic wages and labor market search. Journal of Labor Economics, 16(4), 815–847.

Jacobson, L. S., LaLonde, R. J., & Sullivan, D. G. (1993). Earnings losses of displaced workers. The American Economic Review, (pp. 685–709).

Jardim, E. S., Solon, G., & Vigdor, J. L. (2019). How prevalent is downward rigidity in nominal wages? evidence from payroll records in washington state. Tech. rep., National Bureau of Economic Research.

Jarosch, G. (2021). Searching for job security and the consequences of job loss. Tech. rep., National Bureau of Economic Research.

Johnson, S. G. (2014). The nlopt nonlinear-optimization package.

Kaelo, P., & Ali, M. (2006). Some variants of the controlled random search algorithm for global optimization. Journal of Optimization Theory and Applications, 130(2), 253–264.

Kaur, S. (2019). Nominal wage rigidity in village labor markets. American Economic Review, 109(10), 3585–3616.

Kimball, M. S., Sahm, C. R., & Shapiro, M. D. (2008). Imputing risk tolerance from survey responses. Journal of the American Statistical Association, 103(483), 1028–1038.
URL https://doi.org/10.1198/016214508000000139

Kimball, M. S., Sahm, C. R., & Shapiro, M. D. (2009). Risk preferences in the psid: Individual imputations and family covariation. American Economic Review, 99(2), 363–68.

Krueger, A. B., & Meyer, B. D. (2002). Labor supply effects of social insurance. Handbook of public economics, 4, 2327–2392.

Lachowska, M., Mas, A., Saggio, R., & Woodbury, S. A. (2021). Do workers bargain over wages? a test using dual jobholders. Tech. rep., National Bureau of Economic Research.

Lachowska, M., Mas, A., & Woodbury, S. A. (2020). Sources of displaced workers' long-term earnings losses. American Economic Review, 110(10), 3231–66.

Lamadon, T., Mogstad, M., & Setzler, B. (2019). Imperfect competition, compensating differentials and rent sharing in the us labor market. Tech. rep., National Bureau of Economic Research.

Lentz, R. (2009). Optimal unemployment insurance in an estimated job search model with savings. Review of Economic Dynamics, 12(1), 37–57.

Lise, J. (2013). On-the-job search and precautionary savings. Review of Economic Studies, 80(3), 1086–1113.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. Econometrica, (pp. 995–1026).

Moffitt, R. A. (2002). Welfare programs and labor supply. Handbook of Public Economics, 4, 2393–2430.

Mortensen, D. T. (1977). Unemployment insurance and job search decisions. ILR Review, 30(4), 505–517.

Mortensen, D. T., & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. The Review of Economic Studies, 61(3), 397–415.

Pannenberg, M. (2010). Risk attitudes and reservation wages of unemployed workers: evidence from panel data. Economics Letters, 106(3), 223–226.

Phelps, E. S. (1997). Rewarding work: How to restore participation and self-support to free enterprise. Harvard University Press.

Pinheiro, R., & Visschers, L. (2015). Unemployment risk and wage differentials. Journal of Economic Theory, 157, 397–424.

Pissarides, C. A. (1974). Risk, job search, and income distribution. Journal of Political Economy, 82(6), 1255–1267.

Postel-Vinay, F., & Robin, J.-M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. Econometrica, 70(6), 2295–2350.

Rendon, S. (2006). Job search and asset accumulation under borrowing constraints. International Economic Review, 47(1), 233–263.

Robin, J.-M. (2011). On the dynamics of unemployment and wage distributions. Econometrica, 79(5), 1327–1355.

Rowan, T. H. (1990). Functional stability analysis of numerical algorithms. Ph.D. thesis, The University of Texas at Austin.

Sahin, A., Song, J., & Hobijn, B. (2010). The unemployment gender gap during the 2007 recession. Current Issues in Economics and Finance, 16(2).

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. American Economic Review, 95(1), 25–49.

Solow, R. M. (1990). The labor market as a social institution. Blackwell.

Sullivan, J. X. (2008). Borrowing during unemployment unsecured debt as a safety net. Journal of Human Resources, 43(2), 383–412.

Syverson, C. (2004). Product substitutability and productivity dispersion. Review of Economics and Statistics, 86(2), 534–550.

Syverson, C. (2011). What determines productivity? Journal of Economic literature, 49(2), 326–65.

Taber, C., & Vejlin, R. (2016). Estimation of a roy/search/compensating differential model of the labor market. Tech. rep., National Bureau of Economic Research.

University of Groningen and University of California, Davis (2021). Share of labour compensation in gdp at current national prices for united states [labshpusa156nrug]. Retrieved from FRED, Federal Reserve Bank of St. Louis, Apr. 4, 2021.

Winter-Ebmer, R. (1995). Does layoff risk explain the firm-size wage differential? Applied Economics Letters, 2(7), 211–214.

A Relaxation of modeling assumptions

In this section, I discuss the literature regarding some of the simplifying modeling assumptions that have been made to direct focus on the layoff margin of interest and explore how relaxing those assumptions would play out in the model.

A.1 Allowing for positive shocks

In the main model, following Mortensen & Pissarides (1994) shocks to match productivity can be positive or negative conditional on the current value, yet total productivity can never exceed the initial match value. While it may be of interest to consider an environment where the shocks could also increase the match value, the assumption in the model is without loss of generality from the perspective of layoff risk.

To see why, suppose instead that $F(\varepsilon)$ has both positive and negative support. This enters into the model as a substitution of the lower bound $\underline{\varepsilon}$ where in the primary model the lower bound is 0.

$$\Pi_\varepsilon = \frac{1}{1+\rho}\left(\alpha x - \varepsilon + \lambda_1 P_s(x)\Pi_\varepsilon + \eta \int_{\underline{\varepsilon}}^{\varepsilon^*(x)} \Pi_{\varepsilon'} f(\varepsilon') + (1-\lambda_1-\eta)\Pi_\varepsilon\right) \tag{22}$$

Repeating the steps for the solution from section 3 with this new lower bound, the threshold termination shock becomes

$$\varepsilon^*(x) = \alpha x + \frac{\eta}{D}\int_{\underline{\varepsilon}}^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon. \tag{23}$$

Because a firm will terminate a match only after a negative shock, there is now a subdomain of the support of ε that is irrelevant to that decision, namely everything up to $F(0)$: $\varepsilon^*(x) > 0 \; \forall \; x$. Mathematically, the termination equation can be rewritten

$$\varepsilon^*(x) = \alpha x + \frac{\eta}{D} \int_0^{\varepsilon^*(x)} F(\varepsilon) d\varepsilon + \frac{\eta}{D} F(0). \tag{24}$$

One can see immediately that the dynamics between firm productivity and layoff risk are unchanged.

A.2 Relaxing wage rigidity

The assumption that wages are constant for the duration of a match is common in basic wage posting models in the search literature (e.g. Burdett & Mortensen (1998)), and might be considered an augmented version of a partial job search search model where an exogenous parametric wage offer distribution is assumed as a primitive of the model. The wage posting frame of the model here differs from models with renegotiation after outside offers arrive (Dey & Flinn (2005), Cahuc et al. (2006)), and the premise in Mortensen & Pissarides (1994) that wages are renegotiated following every transitory productivity shock. That assumption of period-by-period renegotiation has been criticized in the macro literature for being unable to capture the volatility of unemployment fluctuations (Shimer, 2005). Proposed alternatives, such as Hall (2005), address the issue by modifying the method of wage determination in the model to deliver some form of rigidity.

More broadly, there is a long documented history that firms do not like to adjust wages in the face of downturns. Azariadis (1975) sought to understand the "normal industrial practice of laying off unneeded workers and paying unchanged wages to the rest of the work force" as opposed to adjusting wages. He proposed that risk-neutral firms act not only as employers but also as insurers for risk-averse workers. More recent evidence that firms act as insurers comes from Guiso et al. (2005), who find that transitory shocks to firms are not passed through to workers. Additional empirical evidence supporting the assumption of wage rigidity includes Altonji & Devereux (1999), who observe that downward wage cuts without changes in job status are rare; Haefke et al. (2013), who show that new hires have much more flexible wages than current workers; and Barattieri et al. (2014), who conclude that workers who switch jobs have a much higher probability of changing wages than workers who continue to be employed at the same firm. (Friedrich et al., 2019) find some pass-through of permanent firm shocks and less for transitory shocks to match productivity. Qualitatively, in a book titled *Why Wages Don't Fall During a Recession*, through interviews with executives and recruiters, Bewley (2009) learns that employers prefer not to cut wages following negative demand shocks because they believe it will reduce worker morale and engagement with the firm's objectives. Similarly, Campbell III & Kamlani (1997) survey 184 firms and document resistance to wage reductions for many reasons, particularly because of employee morale and effort and the asymmetric impact of lowering versus raising wages, replacement costs, and risk of losing their most productive workers. More recently, Jardim et al. (2019) and Elsby & Solon (2019) argue that nominal wage cuts are more prevalent than found in studies using household survey data, finding that about 15-25% of job stayers experience nominal wage cuts in a year. This is contradicted by Grigsby et al. (2021), who use administrative data to show that only 2.5% of job stayers experienced a wage cut between 2008-2016. An additional implication of the wage rigidity in the model is that unemployed workers will underbid current employees by being willing to accept a lower wage for an equally productive match. Solow (1990) believed that the absence of underbidding could explain observed patterns in downward wage rigidity, but Fehr & Falk (1999) found "massive" evidence of just such

underbidding by workers in an experimental setting, yet refusal by firms to accept that underbidding. The explanation offered is that labor markets are characterized by incomplete contracts with fixed wages, and since effort is unobservable, firms prioritize fairness to maintain worker cooperation, which wage reductions put in jeopardy.

In the present paper, the objective is not to explain the phenomenon of wage rigidity, but rather to incorporate this stylized fact so as to support conducting the analysis of interest. The primary model therefore assumes that a match involves a fixed-wage contract, such that wages are not affected by shocks to the match productivity. Suppose instead that we maintained the sharing rule for the match product, yet allowed it to update fully following each ε shock. This is reflected in the firm's value function by making flow profits equal to $\alpha(x-\varepsilon)$ instead of $\alpha x - \varepsilon$

$$\Pi_\varepsilon = \frac{1}{1+\rho}\left(\alpha(x-\varepsilon) + \lambda_1 P_s(x)\Pi_\varepsilon + \eta \int_0^{\varepsilon^*(x)} \Pi_{\varepsilon'} f(\varepsilon') + (1-\lambda_1-\eta)\Pi_\varepsilon\right) \tag{25}$$

In the simplest case, assume that the current draw of ε does not affect the probability that a worker leaves a firm when faced with an outside offer (or affects it in a small enough way that we can assume it is constant for the sake of illustration). Then

$$\begin{aligned} \alpha\varepsilon^*(x) &= \alpha x + \alpha\frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon \\ \implies \varepsilon^*(x) &= x + \frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon. \end{aligned} \tag{26}$$

The mapping from productivity to the maximum shock the firm is willing to absorb is a monotonic transformation from that implied by the primary model. Specifically, the firm is willing to absorb larger shocks, yet the rank ordering of those magnitudes is still determined by x.

Extending this modified approach to the policy experiment of a permanent downturn where $\tilde{x} = x - s$, it matters whether we assume wages can update immediately. Here the wage rigidity assumption plays a larger role. If wages update fully, then the entire wage distribution shifts down such that current workers no longer earn rents greater than they would get by accepting identical jobs in the new labor market. Wages and layoff risk are remapped roughly to $w(\tilde{x}-\varepsilon) = w(x-s-\varepsilon)$ and $\delta(\tilde{x}) = \delta(x-s)$. Every employed worker suddenly finds themselves earning less and facing higher layoff risk with a worse outside option given the new offer distribution. Therefore, we still expect increased job churning, greater unemployment, and reduced lifetime income, yet the wedge disappears between the still employed and the newly unemployed and employed workers.

An in-between approach is to assume imperfect pass-through from shocks to wages at rate $\alpha \leq \beta \leq 1$, as in

$$\Pi_\varepsilon = \frac{1}{1+\rho}\left(\alpha x - \beta\varepsilon + \lambda_1 P_s(x)\Pi_\varepsilon + \eta \int_0^{\varepsilon^*(x)} \Pi_{\varepsilon'} f(\varepsilon') + (1-\lambda_1-\eta)\Pi_\varepsilon\right) \tag{27}$$

Holding D constant, this yields termination condition

$$\beta\varepsilon^*(x) = \alpha x + \beta\frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon$$

$$\implies \varepsilon^*(x) = \frac{\alpha}{\beta}x + \frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon \tag{28}$$

which converges to the main case in the model or the case with full pass-through of shocks as $\beta \to 1$ or $\beta \to \alpha$. One can observe that this still has no practical impact on the baseline results of the model, while reducing but not eliminating the aforementioned wedge in the policy simulation. As a robustness check, after conducting the main policy simulation I repeat the exercise for several value of β to demonstrate the degree to which the wage rigidity assumption influences the main conclusions of the paper.

A.3 Allowing for wage bargaining

Even if wages are negotiated, and renegotiated, the Mortensen & Pissarides (1994) model still results in firms having a productivity threshold below which they will terminate a match, and this threshold is increasing with the initial productivity of the match. That model delivers a tractable bargaining solution by assuming no on-the-job search and homogeneous, risk-neutral workers. Once workers are allowed to differ and to weigh outside offers, which are valued differently over time according to worker preferences and the current wage, the bargaining solution quickly becomes unwieldy. Furthermore, survey evidence reports that two-thirds of workers did not negotiate their current salary (Hall & Krueger, 2012), and empirical estimates of wage bargaining shows that it is uncommon at low and medium wage jobs (Cahuc et al. (2006), Card et al. (2016)), and that the magnitude of the bargaining that does take place is small, with an elasticity with respect to increases in the outside option on the order of just 0.1-0.2% percent (Caldwell & Harmon (2019), Lachowska et al. (2021)). To "flip" the implications in the model to get layoff risk to increase with wages, bargaining power would need to rise so steeply with match productivity as to cause profits to decrease, which does not accord with intuition, the magnitude of the bargaining power estimates in the literature, or the empirical evidence that, if anything, the labor share of income is decreasing with firm productivity (Dorn et al., 2017). Because the bargaining approach does not deliver a fundamentally different conclusion regarding the relationship between productivity and job security, while adding firm and worker heterogeneity and on-the-job search to the model, I make the simplifying assumption that there is no bargaining. Compositional effects could arise if firms could bargain differently with different groups such that more risk averse workers received a lower share of the surplus, yet that assumption would entail making unrealistic or illegal assumptions like that firms perfectly observe risk preferences and asset holdings of workers, and can pay identically productive workers differently. If instead assets and risk preferences are unobserved, firms will bargain based on expectations, again treating all workers the same.

The primary model assumes a sharing rule for the match output rather than Nash Bargaining over the match surplus as is common in the literature. One reason models with Nash Bargaining are able to include it is because of the assumption that workers are risk neutral. This yields the common result that firms and workers split the match surplus. Without that assumption, the Nash Bargaining solution quickly becomes convoluted, although other solution concepts may still be viable. As a simple illustration, assume a much simplified version of the model where there is no on the job search, no saving, and all workers are

homogeneously risk averse. Assume that firms and workers renegotiate wages following every shock to the match value. Denote the offer arrival rate λ.

The value to a worker when unemployed is

$$\rho U = u(b) + \lambda \int_x \max\{V(x,0) - U, 0\}\, d\Gamma(x) \tag{29}$$

and when employed is

$$(\rho + \eta)V(x,\varepsilon) = u(w(x,\varepsilon)) + \eta \int_\varepsilon \max\{V(x,\varepsilon), U\}\, dF(\varepsilon) \tag{30}$$

Assume as before that the firm's outside option is 0, and the value of a filled position is

$$(\rho + \eta)\Pi(x,\varepsilon) = x - \varepsilon - w(x,\varepsilon) + \eta \int_\varepsilon \max\{\Pi(x-\varepsilon), 0\}\, dF(\varepsilon) \tag{31}$$

Denote the relative bargaining power of workers as α. According to the axiomatic Nash Bargaining approach, wages will be chosen to maximize the weighted joint surplus:

$$w(x,\varepsilon) = \arg\max_w \left\{[V(x,\varepsilon) - U]^\alpha [\Pi(x,\varepsilon)]^{1-\alpha}\right\} \tag{32}$$

Dropping the parentheses, the first order condition is

$$\alpha \frac{V'}{V - U} = -(1-\alpha)\frac{\Pi'}{\Pi} \tag{33}$$

which, recognizing that the future terms do not depend on the current wage, can be rewritten as

$$V - U = \alpha(u'(w)\Pi + V - U). \tag{34}$$

Unlike in the typical case, where the worker receives a constant share of the match surplus $\Pi + V - U$, with risk aversion the worker receives a share of the match surplus that is decreasing in the wage. Hence, while wages should be increasing in productivity x, the worker's share of that product will not be rising. More productive firms should earn higher flow profits.

To see that the termination threshold is still increasing in x (layoff risk is decreasing in x) under a bargaining assumption , consider the termination condition where $\Pi(x, \varepsilon^*(x)) = 0$:

$$\varepsilon^*(x) = x - w(x,\varepsilon^*(x)) + \frac{\eta}{\rho + \eta} \int_0^{\varepsilon^*(x)} F(\varepsilon) d\varepsilon \tag{35}$$

Simplifying notation and taking the derivative with respect to x, we get

$$\varepsilon^{*'} = \frac{1}{1 + w' - \frac{\eta}{\rho+\eta} F(\varepsilon^*)} \tag{36}$$

which is positive given that wages are increasing in firm productivity.

A.4 Multiplicative productivity shocks

The primary models follows Mortensen & Pissarides (1994) and others by treating transitive shocks to firm productivity as additive. In theory, they could enter differently, for instance as a multiplier. This implies certain assumptions about productivity shocks, particularly that more productive firms face larger variance shocks, which seems less realistic than the additive implication that shocks are relatively more important to less productive firms. For that reason this case is left unsolved algebraically. To reiterate, the goal of the model is not to prove that a certain wage-risk relationship holds under all possible assumptions and modeling scenarios, but instead that a surprising relationship results given a few basic, reasonable assumptions.

A.5 The firm's outside option

In the main model I assume that there is free entry at all levels of productivity, driving the outside option to zero for all firms. This may not hold if match productivity x is instead an inherent characteristic of a firm that is known before the match is formed and carries over after match destruction. To investigate that scenario, here I write down a firm's value function for maintaining a vacancy and characterize the job termination threshold shock $\varepsilon^*(x)$ when the vacancy value may not be zero.

Suppose the flow cost of maintaining a vacancy is c. Let the value of a vacancy be $F(x)$ where

$$F(x) = \frac{1}{1+\rho}\left(c + \lambda_0 P_{A|x} \Pi(x,0) + \lambda_0(1-P_{A|x})F(x) + (1-\lambda_0)F(x)\right) \tag{37}$$

where $P_{A|x}$ denotes the probability of a worker accepting an offer if matched with the firm. For simplicity, suppose that all jobs are above every worker's reservation wage in expectation, such that we assume $P_{A|x} = 1$. Then

$$(\lambda_0 + \rho)F(x) = c + \lambda_0 \Pi(x,0). \tag{38}$$

The value of a filled vacancy is

$$D\Pi(x,\varepsilon) = \alpha x - \varepsilon + \frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon')d\varepsilon' + \lambda_1 \overline{P_s} F(x) \tag{39}$$

such that at the termination threshold

$$(\rho + \eta + \delta_0)\Pi(x,\varepsilon^*) = \alpha x - \varepsilon^* + \frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon')d\varepsilon' \tag{40}$$

because

$$D = \rho + \lambda_1 \overline{P_s}(x) + \eta + \delta_0 \tag{41}$$

and we know from the model that at the termination threshold, $F(x) = \Pi(x,\varepsilon^*)$.

Substituting, we get

$$\frac{\rho+\eta+\delta_0}{\lambda_0+\rho}c + \frac{(\rho+\eta+\delta_0)\lambda_0}{\lambda_0+\rho}\Pi(x,0) = \alpha x - \varepsilon^* + \frac{\eta}{D}\int_0^{\varepsilon^*(x)} F(\varepsilon')d\varepsilon' \tag{42}$$

or

$$\varepsilon^*(x) = \alpha x + \frac{\eta}{D} \int_0^{\varepsilon^*(x)} F(\varepsilon')d\varepsilon' - \frac{\rho + \eta + \delta_0}{\lambda_0 + \rho}(c + \lambda_0 \Pi(x, 0)) \tag{43}$$

This looks similar to the termination in the main model, which is increasing in x, plus two additional terms. The vacancy cost is a negative number, meaning the costlier a vacancy, the larger a shock the firm will absorb before terminating a match. Although I have assumed that c is the same for all firms, it is easy to imagine that in a richer model, a vacancy is costlier for a more productive position, supposing the Human Resources Department spends more to find a new manager than it does for a new mail clerk. Counterbalancing this dynamic is the value of a newly filled position, which is increasing in x. The higher the probability that the firm can quickly find a new worker, the more willing it is to destroy a match and try again.

B Estimation results under alternative decisions

In this appendix, I present the results from estimating the model given different sampling and estimation decisions to show how that affects the estimated structural parameters. First, Table 15 show how the estimated parameters and function value change as the number of discrete points used to approximate the distribution of risk preferences varies. The parameters are generally consistent across versions of the estimator, and the objective function is lowest at the main specification in the paper with 5 discrete points approximating the continuous distribution.

Table 15: Parameter estimates for different numbers of discrete points

Parameter	3 points	4 points	5 points	6 points	7 points
λ_0	0.183	0.204	0.186	0.290	0.197
λ_1	0.089	0.102	0.086	0.125	0.087
α_x	2.461	2.615	3.283	1.950	3.494
β_x	7.714	8.532	9.311	7.361	9.874
θ_ε (x100)	0.061	0.069	0.063	0.075	0.068
δ_0 (x100)	0.153	0.193	0.164	0.263	0.246
η	0.040	0.056	0.051	0.161	0.046
α_γ	2.789	2.927	3.028	4.377	4.632
β_γ	4.368	5.064	4.700	9.423	10.397
b	604.974	604.974	604.974	604.974	604.974
fval	1981.235	1944.540	1813.097	2346.095	2074.417

Although the residualized measures for assets and wages correspond more closely with the conceptual variation in the model, one might be curious to know how influential that adjustment is relative to estimation on the raw data. Table 16 contains the results from estimating the model without adjusting assets or wages for variation due to observable worker heterogeneity or time. The shock arrival rates, offer distributions, and risk preference distribution vary somewhat yet are all generally in line with those estimated for the residualized model, such that using this version would not meaningfully alter the policy implications calculated.

Table 16: Parameter estimates using nonresidualized data

Parameter	Estimate
λ_0	0.209
λ_1	0.119
α_x	1.687
β_x	8.411
θ_ε (x100)	0.081
δ_0 (x100)	0.228
η	0.104
α_γ	4.701
β_γ	9.678
b	604.970
$\min(x)$	1899.801
$\max(x)$	21246.303
fval	1841.917
Calibrated	
α	0.4
ρ	0.006
r	0.00125

C Validation of estimator performance

To verify that the estimator is coded correctly and can recover the structural parameters of the model, I generated a sample of 20,000 labor market histories governed by a calibrated set of parameters and then ran the estimator on that sample. The main specification of the model allows for heterogeneity in risk preferences governed by a distribution known up to the parameter values. Table 17 shows that the procedure of sorting initial asset values and assigned risk preferences draws allow the estimator to recover the true structural parameters of the model. For example, the true mean coefficient of relative risk aversion is 0.86 and the estimated mean is 0.91, and the 0-1 normalized mean of the beta distribution for productivity offers is estimated at 0.222 with a standard deviation of 0.132 compared to true values of 0.200 and 0.137.

Table 17: Validation of estimator against known parameters - heterogeneous workers

Parameter	Set Value	Estimate
λ_0	0.25	0.210
λ_1	0.15	0.172
α_x	1.50	1.988
β_x	6.00	6.962
θ_ε (x100)	0.05	0.050
δ_0 (x100)	0.25	0.210
η	0.03	0.036
μ_γ	2.00	1.997
σ_γ	5.00	4.587
b	300.00	300.000
$\min(x)$	1400.00	1410.014
$\max(x)$	22000.00	17070.048
fval	1006.50	139.873
Calibrated		
α	0.4	
ρ	0.006	
r	0.00125	

D Hypothetical risk gambles in the PSID

One empirical approach to estimating risk preferences is to elicit responses to survey questions regarding hypothetical gambles over income. For the PSID, this was done in the 1996 wave, where employed respondents were asked if they would be willing to trade a guaranteed job for life for one that could either double their income or reduce it by X%. Respondents are classified into one of six groups depending on where they drew the line. Expected utility theory provides an upper and lower bound for the preferences that are consistent with the income gamble that each participant accepts. Kimball et al. (2009) have impute estimates of risk tolerance preference parameters for each group, accounting for measurement error in responses, with the imputed coefficient of relative risk aversion ranging from 1.4 to 6.7. They fit a lognormal distribution to the data and calculate a mean of 3.8 with the vast majority of workers having a coefficient above 2.2. In my model, risk preferences this steep would imply that many workers want to many multiples of annual income in precautionary savings. Table 18 reports the full distribution of risk groups and the corresponding Kimball et al. (2009) imputed coefficients of relative risk aversion in the 1996 wave of the PSID.

Table 18: The Distribution of Risk Aversion in the PSID (Kimball et al., 2009)

Group	γ	N	%
1	6.7	1,230	30.15
2	4.2	728	17.85
3	3.5	629	15.42
4	2.8	620	15.20
5	2.2	582	14.27
6	1.4	290	7.11
	Total	4,079	100.00

While it is nonetheless tempting to leverage these stated preferences in some way to identify risk aversion, I choose not to because they do not appear to predict savings behavior. I demonstrate this in Figures 21 and 22, which use the 2003 wave as it is the closest sample year to when preferences were elicited in 1996. Figure 21 plots local polynomial regressions of savings relative to earnings separately for the most risk averse and least risk averse PSID respondents based on the hypothetical gamble questions. There is no discernible difference in their savings behavior. The same puzzle appears in Figure 22, where I plot the density of savings to earnings ratios for the purportedly most and least risk averse groups. Once again, the two groups exhibit remarkably similar behavior. Because of this lack of expected differences in behavior, I avoid using these hypothetical gambles to identify the model here, leaving further exploration of the empirical content of the risk preference survey questions to future work.

Figure 21: Savings per earnings by risk preference level (PSID)

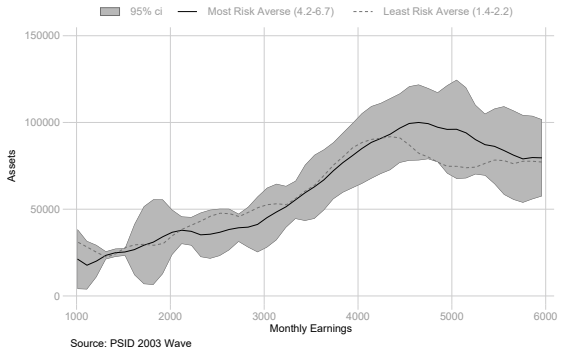

Figure 22: Savings-income ratios by risk preference level (PSID)

E Value function details

E.1 Value function for the firm

Recall that the present discounted expected value of a match for the firm be

$$\Pi_\varepsilon = \frac{1}{1+\rho}\left(\alpha x - \varepsilon + \lambda_1 P_s(x)\Pi_\varepsilon + \eta \int_0^{\varepsilon^*(x)} \Pi_{\varepsilon'} f(\varepsilon') + (1-\lambda_1-\eta-\delta_0)\Pi_\varepsilon + \delta_0 0\right) \tag{44}$$

where P_s is the probability that a worker stays after receiving an outside offer, the discount factor $\beta = \frac{1}{1+\rho}$, $f(\varepsilon)$ is the density of ε shocks, and $\varepsilon^*(x)$ is the threshold for a shock that is large enough for the firm to terminate the match. In a new match, Π_0, $\varepsilon = 0$.

Rewriting and collecting terms,

$$(\rho + \lambda_1 \overline{P_s}(x) + \eta + \delta_0)\Pi_\varepsilon = \alpha x - \varepsilon + \eta \int_0^{\varepsilon^*(x)} \Pi_{\varepsilon'} f(\varepsilon') \tag{45}$$

Integrating by parts on the righthand side, the value can be rewritten as

$$(\rho + \lambda_1 \overline{P_s}(x) + \eta + \delta_0)\Pi_\varepsilon = \alpha x - \varepsilon - \eta \int_0^{\varepsilon^*(x)} \Pi'_{\varepsilon'} F(\varepsilon')d\varepsilon' \tag{46}$$

Further, we can see from equation 45 that

$$\Pi'_\varepsilon = \frac{-1}{\rho + \lambda_1 \overline{P_s}(x) + \eta + \delta_0}. \tag{47}$$

For ease of notation, define that denominator term as

$$D = \rho + \lambda_1 \overline{P_s}(x) + \eta + \delta_0. \tag{48}$$

Then, substituting back in to equation 46,

$$D\Pi_\varepsilon = \alpha x - \varepsilon + \frac{\eta}{D} \int_0^{\varepsilon^*(x)} F(\varepsilon')d\varepsilon'. \tag{49}$$

E.2 Layoff risk is decreasing in firm productivity

This can be seen by taking the derivative of the termination condition with respect to x:

$$\varepsilon^{*'}(x) = \alpha + \frac{\eta}{D}\left(\varepsilon^{*'}(x)F(\varepsilon^*(x))\right) + -\frac{\eta}{D^2}D'(x)\left(\int_0^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon\right)$$

$$\left(1 - \frac{\eta}{D}F(\varepsilon^*(x))\right)\varepsilon^{*'}(x) = \alpha - \frac{\eta}{D^2}\left(-\lambda_1 P'_s(x)\right)\int_0^{\varepsilon^*(x)} F(\varepsilon)d\varepsilon$$

$$\left(1 - \frac{\eta}{D}F(\varepsilon^*(x))\right)\varepsilon^{*'}(x) = \alpha - \frac{\eta}{D^2}\left(-\lambda_1 P'_s(x)\right)\left(\varepsilon^*(x) - \alpha x\right)\frac{D}{\eta}$$

$$\varepsilon^{*'}(x) = \frac{\alpha + \lambda_1 P'_s(x)(\varepsilon^*(x) - \alpha x)/D}{1 - \frac{\eta}{D}F(\varepsilon^*(x))}$$

The numerator is positive so long as P'_s is not "very negative," which is to say we assume any marginal increase in layoff risk as productivity rises is not drastic relative to the utility gain from the marginal increase in wage. This must always be true if we assume that in equilibrium a more productive firm offers no worse value for the most risk averse worker than does a less productive firm, because then $P'_s(x) \geq 0$. Since $0 < \eta/D < 1$ and $0 \leq F(\varepsilon^*(x) \leq 1$, the entire expression is therefore positive.

E.3 Worker indifference condition and risk

Unemployed workers do not value layoff risk when deciding whether to accept a job. To see this, substitute for the discount factor β with the discount rate ρ as in the firm section, where $\beta = \frac{1}{1+\rho}$. We can rewrite the

value functions

$$\rho U(a) = u(c) + U(a') - U(a) + \lambda_0 \int \max\{V(a', x) - U(a'), 0\}\, d\Gamma(x) \tag{50}$$

$$\rho V(a, x) = u(c) + V(a', x) - V(a, x) + \lambda_1 \int \max\{V(a', x') - V(a', x), 0\} + \delta\left(U(a') - V(a', x)\right) \tag{51}$$

A worker is indifferent between accepting or rejecting a job offer when $V(a, x') = V(a, x)$. Plugging into the value function for employment at that indifference point yields the condition

$$\begin{aligned} u(c(a, w')) + \delta'\left(U(a) - V(a, x)\right) &= u(c(a, w)) + \delta\left(U(a) - V(a, x)\right) \\ \implies u(c(a, w')) &= u(c(a, w)) + (\delta' - \delta)(V(a, w) - U(a)) \end{aligned} \tag{52}$$

The indifference condition for an unemployed worker with a certain level of savings is $U(a) = V(a, x)$. Given assets and risk aversion, the reservation wage and risk, $r(w, \delta; a, \gamma)$, can be characterized by considering the special indifference case where the value of employment equals the value of unemployment. Making this substitution into the value functions yields

$$\begin{aligned} u(c(a, b)) + \lambda_0 \int \max\{V(a', x) - U(a'), 0\}\, d\Gamma(x) &= u(c(a, r)) + \lambda_1 \int \max\{V(a', x) - U(a'), 0\}\, d\Gamma(x) \\ \implies u(c(a, r)) &= u(c(a, b)) + (\lambda_0 - \lambda_1) \int \max\{V(a', x) - U(a'), 0\}\, d\Gamma(x) \end{aligned} \tag{53}$$

The righthand side does not depend on the reservation risk level. Suppose there were two reservation pairs (r_1, δ_1) and (r_2, δ_2) where $r_1 \neq r_2$. The only way to equate $u(c(a, r_1))$ and $u(c(a, r_2))$ is if $a_1' \neq a_2'$, which is a contradiction given the indifference condition assumes equal assets. Given assets and risk preferences, the reservation wage is independent of the risk offer. This creates a tension regarding risk preferences, as more risk averse workers have lower reservation wages, which come with a higher probability of sending them back into unemployment.

F CPS details and subgroup analysis

This appendix describes the construction of the CPS sample used for the motivating figures in section 1 and presents additional figures for various demographic subgroups and reported reasons for jobs ending to demonstrate that the central pattern is robust to these refinements.

F.1 Sample construction

The data includes every month in 2006, 2008, 2019, and 2020, downloaded from the IPUMS CPS service Flood et al. (2018) and makes use of their constructed user id to track respondents from month to month. The sample is restricted to labor force participants ages 18 to 69 who are not self-employed. Employment status is measured each month of the respondent's participation cycle, and earnings for each respondent are drawn from the most recent month when they reported earnings. All figures are weighted to using the final person-level weights.

F.2 Subgroup analysis

One possible explanation for this pattern is composition: lower earners could be different from higher earners in a way that is systematically related to their likelihood of becoming unemployed (Winter-Ebmer, 1995). On the basis of observable heterogeneity, this does not appear to be the case. As shown in Figure 23, there are no clear differences comparing high school or college educated workers, men and women, black and white workers, workers over and under 40, nor within a specific industry or occupational group; the broad pattern in Figure 1 persists. To reiterate the earlier point, the claim is not that there is no sorting in the labor market; it is that in aggregate as well as conditional on any sorting factors, layoff risk still decreases as earnings rise.

(a) High school vs. college

(b) Male vs. female

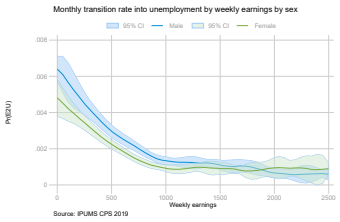

(c) Black vs. white

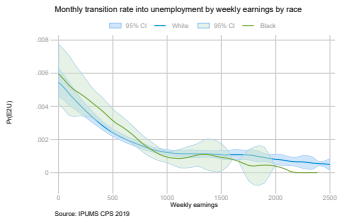

(d) Age under 40 vs. age over 40

(e) Construction sector

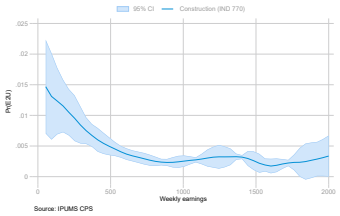

(f) Sales occupations

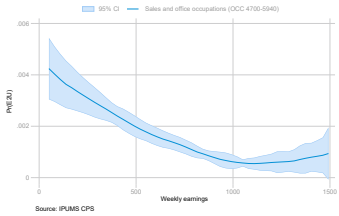

Figure 23: Monthly transitions into unemployment by earnings for various subgroups

A second question involves the distinction between quits and layoffs. It may be that lower earners are more likely to quit voluntarily. Here the charts display all transitions into unemployment because stated reasons for leaving a job may not necessarily reflect the true cause. For instance, a worker may "quit" in anticipation of being laid off. Nonetheless, to show that quit behavior is not driving the observed trend, in Figure 24 I plot separately the transition rate into unemployment across the earnings distribution for people who specifically state that they were laid off versus those who specifically state that they quit their job. Stated quits into unemployment are rare in the CPS data, and the pattern of a hazard declining in earnings is clear when solely examining stated layoffs.

Figure 24: Monthly transition rate into unemployment by weekly earnings - quits vs. layoffs