An Equilibrium Model of Wage and Hours Determination: Labor Market Regulation in the Retail Sector*

Working Paper

Most Recent Version Here

NICK FRAZIER

Rice University

July 2018

ABSTRACT: A recent push to limit the discretion large retailers have over their employee's schedules aims to increase predictability but at an unmeasured cost to market efficiency. Understanding the implications such policies have for a labor market where hours vary and employees have limited control over their own scheduling requires a model of why such jobs exist and also why individuals accept them. This paper formulates and estimates an equilibrium search model with firm and worker heterogeneity that endogenously generates labor contracts, hiring decisions, and search behavior that matches observed patterns in wages, hours, and employment for the U.S. retail sector. I use a novel approach to separately identify the primitives of the supply and demand side optimization problems that incorporates a mixture of stated preference data collected from workers, data on equilibrium retail jobs, and data on employment flows. The empirical results indicate a counterfactual policy that restricts the extent to which hours may vary in a given week does reduce average variability but also results in a two percent decline in aggregate production and leaves most workers worse off in equilibrium.

*I am grateful to Flávio Cunha, Ken Wolpin, Xun Tang, and Russell Cooper for helpful comments and suggestions.

I. Introduction

Technological advances offer firms increasingly finer control over employee scheduling. An algorithm sold by Kronos Incorporated assimilates multiple factors including historical sales, holidays, events, local competition, and weather to provide on-demand optimization of employee scheduling up to the fifteen minute mark. The ability to fine-tune inputs may lead to more efficient firms but comes at a cost to the workers whose schedules are fine-tuned. The extreme example is on-demand scheduling where employers ask employees to maintain availability for a shift but decide later, sometimes hours before, whether or not they will actually work. A large public backlash to on-demand scheduling led U.S. firms to largely abandon the practice (see Nassauer (2016), White (2015), Tabuchi (2015)) and its decline suggests some limits to the gig economy in some industries. However, the larger issue of control over scheduling continues to inspire regulation of the retail and food service sectors including recent measures in San Francisco, Seattle, and New York City.[1] For example, in 2017 the City of Seattle added a requirement that to reduce an employee's hours with one week or less of advance notice, a retail firm has to pay half that reduction back to the employee. These sort of policies directly impact the determination of wages, hours, and employment within the retail industry which in turn accounts for around ten percent of the U.S. labor force and as much as seventeen percent of workers under the age of twenty-five.

To illustrate the extent of variability in hours I consider a sample of jobs drawn from the National Longitudinal Survey of Youth 1997 Cohort (NLSY97) which in 2011 (Round 15) began asking respondents not only what their typical hours per week were, but also what were the most and fewest hours they worked over the last four weeks.[2] Using this survey, Lambert et al. (2014) estimate that around three-quarters of the U.S. workforce in their early thirties experience variability in week-to-week hours. Figure 1 highlights this variability for workers in the retail sector. The solid line plots the median "usual hours per week" of jobs. Around each bin of usual hours are percentiles for the most and fewest hours. The (nearly coincident) dashed line plots median hours for jobs in the same industry but where hours do not vary. These jobs appear identical when summarized by median hours but can imply substantially different sequences of hours and compensation for workers. Looking at the reported range in weekly hours reveals a median of over eight hours—the equivalent of a full shift for retail employees. This variability occurs in an industry where around eighty percent of employees have no direct control over when they work.

The case for regulation often cites survey evidence describing the negative consequences of variability in

[1]The *Fair Workweek Initiative* is a prominent umbrella group that coordinates efforts. See more at http://www.fairworkweek.org

[2]The NLSY97 is a nationally representative sample of approximately 9,000 individuals born between 1980 and 1985. The Work Schedule questions are based on similar ones used in the Workplace and Employee Survey managed by Statistics Canada and the University of Chicago's Work Scheduling Study.

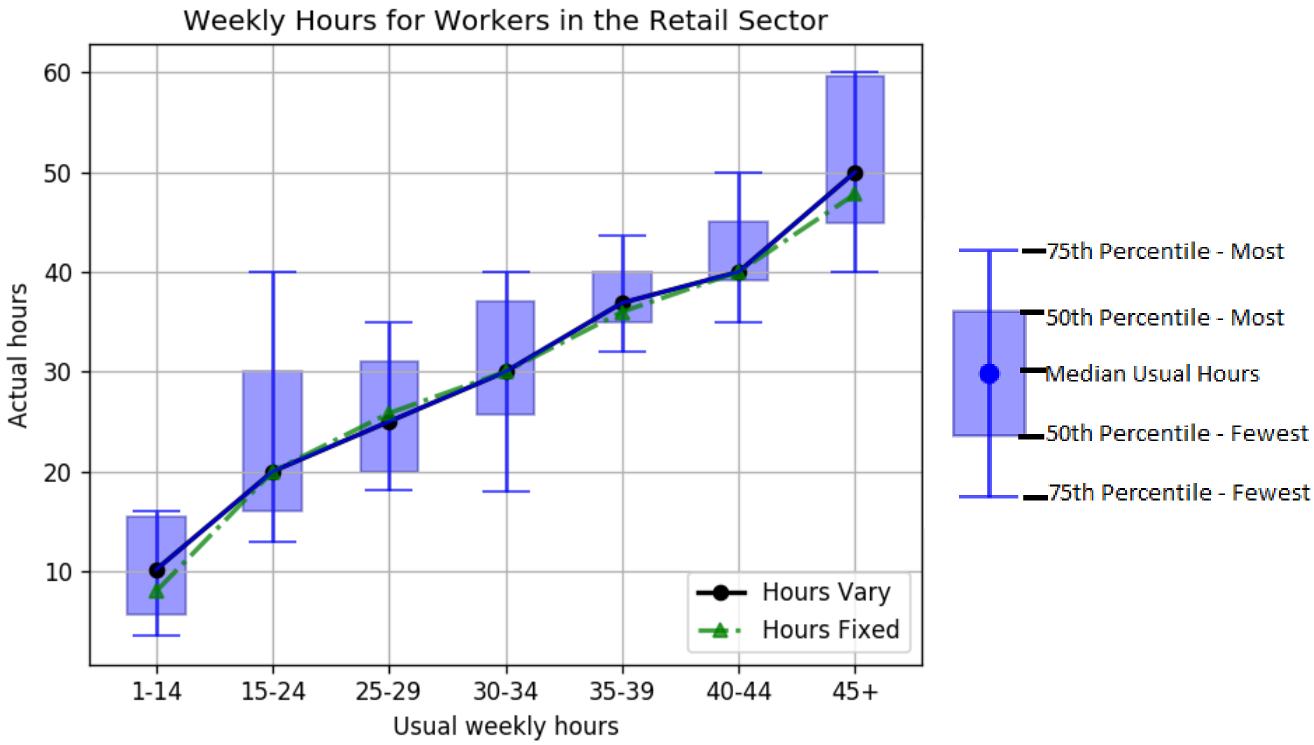

Figure 1: The distribution of actual hours conditional on usual hours for 405 retail employees from the 2013 round of the National Longitudinal Survey of Youth 1997 weighted to be nationally representative. Section V provides additional information about the sample. Based on a similar figure in Lambert et al. (2014) for all industries and reproduced for comparison in Figure 17.

earnings, hours scheduling, and employment instability on employee work-life balance.[3] The evidence on the importance of these job attributes from experimental data is more mixed. Eriksson and Kristensen (2014) use the stated preferences of Danish employees from a discrete choice experiment and find a willingness to pay for flexibility in hours of around twelve percent of the average net wage. Mas and Pallais (2016) embed their discrete choice experiment within the actual hiring process for a call center and find substantial heterogeneity but small average valuations for flexibility and advanced notice of schedule. Some limited but growing evidence on the effect of policy comes from randomized control trials as in Bloom et al. (2015) who find an increase in both productivity and work satisfaction when randomizing the opportunity for Chinese call center employees to work from home.[4] While important, without a model of how these markets work this experimental approach does not fully capture how and through what channels welfare changes. Equilibrium outcomes like sorting, search frictions, and changes in workforce composition, complicate

[3] Reviews of survey evidence are Lambert et al. (2014) and Golden (2015). Examples of evidence from surveys include Golden et al. (2013), Henly and Lambert (2014), Ingre et al. (2012), Lambert (2008), Lambert et al. (2012), Reynolds and Aletraris (2006), Reynolds (2003), and Stewart and Swaffield (1997).

[4] See also the relevant study underway described in Lambert et al. (2015).

attempts to estimate the valuations of non-wage attributes using observed jobs. Brown (1983) outlines the issues when using hedonic models for this purpose including that estimated valuations may have the wrong sign. Heckman et al. (2010) characterizes the scope of non-parametric identification of hedonic models for a single non-monetary attribute while noting that to separate the buyer's willingness to pay from the seller's willingness to accept requires specifying the underlying structural parameters of buyers and sellers. Fully disentangling the costs and benefits of policy that simultaneously reduces both firm efficiency and the uncertainty employees face over hours and compensation requires a structural model of the labor market that endogenously prices this control over hours.

Given the above evidence on varying weekly hours, one important consideration is how to model the source of volatility. A standard neoclassical labor supply model has employees choosing hours free of constraints from the firm. Although this approach can generate schedules where hours vary or remain the same across weeks, the implication that observed labor hours are then optimally chosen labor hours implies that any restrictions on how hours are set cannot be welfare enhancing. Alternatively, incorporating search frictions provides a way to generate sub-optimal contracts by imposing demand side restrictions on hours. The prototypical search model assumes individuals receive some number of job offers that consist of both a wage and a number of hours drawn from some exogenous joint distribution. Examples include van Soest et al. (1990), Blau (1991), Tummers and Woittiez (1991), Dickens and Lundberg (1993), Bloemen (2000), Bloemen (2008), and Cunha and Frazier (2016). In this framework, individuals choose between jobs that vary in offered compensation and hours but the friction resulting from not being able to freely choose hours leaves room for policy that increases welfare by changing the set of available contracts.

To understand the effect of regulations that restrict the contracting space I formulate an equilibrium search model of hours and wage determination in the retail labor market. In the model, a job offer is a wage and a distribution of hours. The labor market has two distinct sectors which differ in how employers choose hours. In one, the employer can adjust hours in response to a productivity shock so that from the worker's perspective total compensation and hours vary randomly from week to week. In the other, hours are fixed and jobs provide a degenerate distribution of weekly hours and compensation. For convenience I denote the former as the *variable contract sector* and the latter as the *fixed contract sector*. In both cases, wages are set before the shock-contingent marginal productivity of labor is known. These contracts occur between risk averse workers who are heterogeneous in how much they dislike variance in consumption and leisure and risk neutral firms who are heterogeneous in the variance of their shocks to productivity. A high variance firm prefers to adjust hours and match with a less risk averse worker. Worker preferences exhibit constant relative risk aversion and are characterized by heterogeneity in the coefficients of risk aversion for consumption and leisure which in turn have separate marginal distributions in the population. Firms produce using a Cobb-Douglas technology that converts labor hours into the economy's consumption good subject to a firm-specific distribution of productivity shocks. Search frictions in the economy generate imperfect sorting between workers and firm and cause unemployment in equilibrium. Though expected payoffs must satisfy participation constraints, actual hours worked are chosen by the employer given the wage and therefore not necessarily optimal from the perspective of the worker preferring to smooth consumption and leisure.

The model includes multiple channels for policy to affect outcomes and welfare. Workers direct their search towards a specific labor market sector which allows for sorting between worker and firm types across and within sectors. Jobs are matched worker-firm pairs where two-sided worker and firm heterogeneity and search frictions create cross-sectional variation in the composition of observed jobs resulting from equilibrium labor contracts. The bargaining environment endogenously sets the offered wage while accounting for anticipated variation in hours, worker preferences, outside options, and aggregate market conditions. The model allows for endogenous firm entry and hiring decisions to capture how changes in market conditions alter employment levels and equilibrium worker flows. Changes in policy operate through these mechanisms which allows equilibrium adjustments in wages, hours, and employment. Quantifying the welfare impact of policy for individual agents requires disentangling how the primitives of the firm, worker, and aggregate economy determine equilibrium outcomes. For example, low degrees of risk aversion in workers or large returns to adjustment for firms could equally result in high variability in hours with contrary implications for policy. Similarly, a relative scarcity of vacancies offering fixed hours could reduce equilibrium wages and thus overstate the apparent value of predictability to workers.

My approach to separately identifying the primitives relating to worker preferences, firm technology, and matching dynamics combines data from the retail sector on observed jobs and employment flows with data collected from workers containing their stated preferences over the kinds of contracts considered here. I use stated preferences data as a solution to the difficulty of identifying an equilibrium search model without firm data, as discussed in Flinn and Heckman (1982) and especially Flinn (2006) who describes it as a "a nearly impossible task." I elicit preferences using a purpose-made survey instrument that asks respondents to repeatedly choose their most preferred option between two hypothetical job offers and a monetary unemployment benefit. As in the model, job offers contain a wage and distribution of weekly hours which provides a way to trace out preferences over the entire contract space—including non-equilibrium combinations. Wage rates are tailored to the individual but vary randomly across offers for which all attributes are observed. Variation in observed choices allows the econometrician to distinguish and recover common preferences as well as heterogeneity across individuals.

To identify the primitives of firm and matching technology I use data from the retail labor market in equilibrium. My data contains the wage and range of weekly hours for four hundred and five retail sector jobs taken from the nationally representative NLSY97. Around fifty-five percent of these jobs report hours that vary week to week and department stores, supermarkets, and specialty stores comprise the majority of employers. Within the context of the model, variation in hours for a given wage provides information about the parameters of firm technology and the distribution of that firm's productivity shocks. The third source of data is the Job Openings and Labor Turnover Survey from the Bureau of Labor Statistics which provides monthly measures of job openings, separations, and hiring rates specific to the retail sector. These moments combined with data on unemployment and labor market tightness pin down employment levels and flows in the model's equilibrium which identifies the parameters that govern search behavior.

I estimate the model in two stages where I first recover worker primitives from the stated preferences survey and then take those as given when estimating firm and search technology. In the first stage, I estimate a random utility version of the preferences in the model that rationalizes the observed sequences of stated preferences over hypothetical jobs offered in the survey. I use simulated maximum likelihood to estimate the common elements in preferences as well as integrate out and recover the heterogeneity across workers. The resulting coefficients of relative risk aversion for consumption vary between 0.66 and 1.04 with a median value 0.82. I find a larger degree in risk aversion over variation in labor hours where estimated coefficients range from 1.12 to 2.83 and imply Frisch elasticities between 0.35 and 0.89 which align with values found in the micro literature. Notably, these estimates come from the stated preferences of workers within a discrete choice experiment where jobs entail exogenously determined variation in weekly hours. This approach contrasts to the standard approach of applying the neoclassical labor supply model to data with observed hours. I provide suggestive evidence that part of this variation is driven by demographic differences between workers.

The second stage of estimation involves solving for the steady state model equilibrium and simulating data to generate a cross-section of jobs and aggregate worker flows. Conceptually, the simulated data is directly comparable to the equilibrium jobs and employment measures I observe in my data from the retail sector. The parameters underlying firm and search technology are then chosen through a simulated method of moments estimator to match the model equilibrium to that observed in the data. I find an output elasticity with respect to hours of 0.66 which given the Cobb-Douglas production technology implies that labor costs amount to sixty-six percent of firm revenue. This value matches the finding of Cooper et al. (2015) when applying a similar production technology with stochastic productivity to manufacturing data. My model is able to broadly match the observed patterns in wage, hours, and employment.

Given the structural parameters, I conduct a counterfactual policy experiment that limits an employee's hours in a given week to be within eight hours of that employee's mean hours. This restriction reduces the ability of firms to adjust to productivity shocks and has implications for firm entry, hiring behavior, wages, hours, and employment levels. To measure welfare impact I simulate the new equilibrium and calculate the annualized change in annual compensation that makes each worker indifferent between the original and counterfactual equilibrium. The restriction binds for eighteen percent of workers in the new equilibrium and while it increases the proportion of jobs where hours vary, it does reduce the overall average variability. This relatively light constraint reduces aggregate production by two percent and produces a small decline in average worker welfare. However, I find the policy has an uneven effect where some workers end up worse off by the equivalent of two weeks of pre-tax income.

This paper contributes to a literature on estimating equilibrium search and bargaining models with heterogeneous agents that originates with Eckstein and Wolpin (1990) and is summarized in Mortensen and Pissarides (1999). My model introduces job offers that contain a wage and demand side restrictions on hours as well as a multi-sector labor market that allows for directed search and sorting. In terms of econometric approach this paper combines the simulated method of moment approach used in Postel-Vinay and Robin (2002) and Cooper et al. (2007) with an initial stage that recovers the heterogeneous preferences of workers then used as inputs when simulating the equilibrium in the second stage. My model uses the bargaining process developed in Dey and Flinn (2005) and Postel-Vinay and Robin (2002) though I make the simplifying assumption that the worker has zero bargaining power in negotiations. Incorporating data from firms provides a way to relax this assumption and still maintain identification as shown with on-the-job search in Postel-Vinay and Robin (2002) and through a worker bargaining weight as in Flinn (2006). However, achieving identification of an equilibrium search model using only data from workers, or from one side of the market, is itself notable within the literature.

Perhaps the closest models to this paper are Dey and Flinn (2005) and Cooper et al. (2007). In Dey and Flinn heterogeneous workers have preferences over job contracts that include a wage and employer-provided health insurance. The firm's cost of provision is exogenously drawn and risk-neutral workers differ in how much they value having health insurance. They estimate worker and firm primitives using supply side data on equilibrium jobs and employment transitions as well as labor's share of firm revenue which pins down the bargaining power of workers. In contrast, by recovering worker preferences separately, my approach allows variation in observed job contracts to identify precisely how the choice in hours affects firm profitability and productivity without the incorporation of demand side data. In Cooper et al. homogeneous but risk averse workers face uncertainty over compensation and hours but contracts perfectly insure workers against firm- and aggregate-shocks by exactly satisfying their *ex-post* participation constraint. They calibrate primitives related to worker preferences and use a mixture of plant-level data and worker flows to estimate a set of firm and search primitives similar in value to those reported here.

The inclusion of hypothetical data in estimating structural search models has been explored in Bloemen

(2008) who compares a stated preference over hours to observed hours and van der Klaauw (2012) who finds an increase in estimation efficiency when incorporating the subjective expectations of future occupation into a model of occupational choice. My paper introduces its use to the equilibrium search literature as a way to expand the scope of the labor contracts and policies that can be evaluated. My use of discrete choice experiments to recover worker preferences resembles the approach in Eriksson and Kristensen (2014) and Mas and Pallais (2016) who use the exogenous variation in the offered wage to recover distributional estimates of employee willingness to pay for non-monetary job attributes.

The plan of the paper is as follows. In Section II, I formulate my equilibrium model of the labor market. Section IV discusses identification given data and outlines my approach to estimation. In Section V, I describe the sample of retail jobs motivating the policy experiment. Section VI describes the data, methodology, and results of estimating worker preferences and heterogeneity. Section VII contains the analogous content for firm and search technology within the general equilibrium model. Section VIII uses the estimated structural parameters to characterize the welfare costs of various policies and Section IX concludes.

II. Model

A unit measure of infinitely-lived workers and a large but finite measure of firms participate in a labor market consisting of two independent sectors. Importantly, each firm and each worker chooses a specific sector in which to look for an employee or employment. Labor contracts in the *fixed sector* specify a wage and a number of hours set constant across production periods. Contracts in the *variable sector* specify a set wage but permit the firm to adjust hours each period. Consequently, the aggregated labor market contains jobs with the same hours each week and others where hours vary across weeks. The sectors operate separately but concurrently with the stipulation that any contract signed in a sector must match the sector type.

Time is discrete and divided into four phases within each period as described here and depicted in Figure 2. In the first phase, firms decide whether to enter and all unmatched agents search and meet through a frictional matching process described below. In the second phase, firm-worker pairs negotiate contracts that specify a wage and the distribution of potential labor hours. In the third phase firms see their idiosyncratic productivity shock and choose labor hours given the terms of the labor contract. In the fourth and final phase, some matches are exogenously destroyed. All agents discount future periods at the same rate β. The remainder of this section develops the model by describing the segmented structure of the labor market, the primitives of the firm's and worker's problem, and the details and dynamics of the successive phases. The final subsection defines a steady state equilibrium and derives properties of the model.

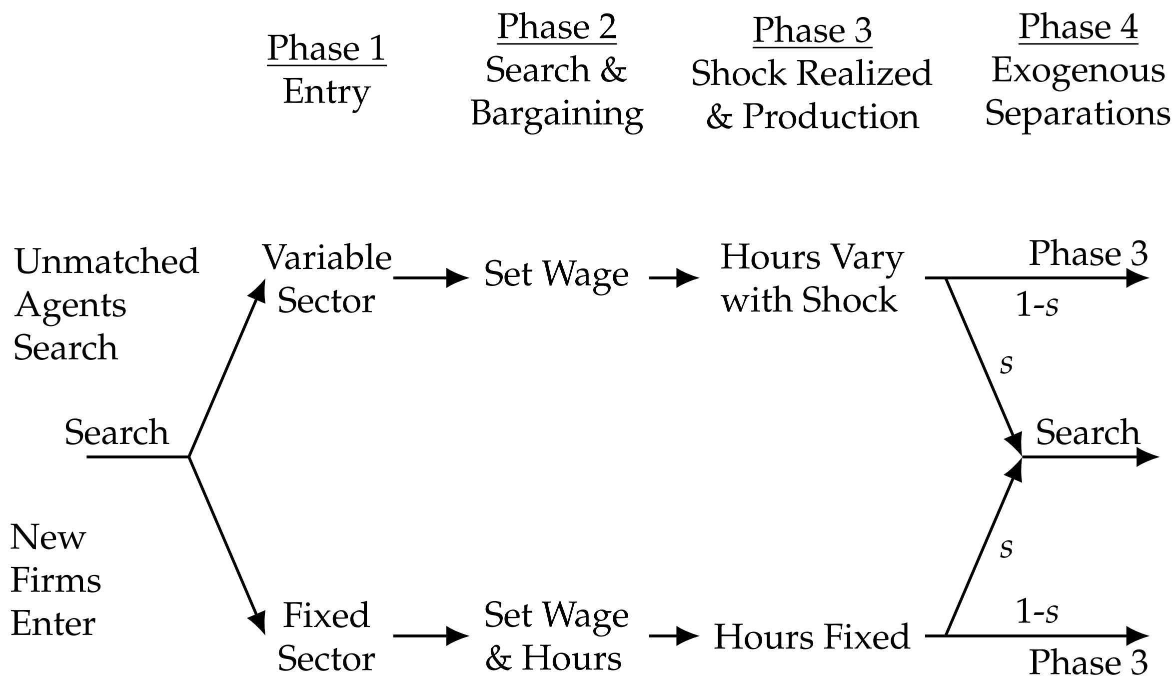

Figure 2: An illustration of the timing in the model. Each discrete period consists of four phases where agents sequentially search, bargain, and produce. Matches are destroyed exogenously at the end of each period with probability s.

A. Firms

Firms own exclusive rights to a technology that converts labor hours into the economy's sole consumption good. A firm may operate a single production site employing a single worker. Firms operate independently and production is a function of the worker's labor hours and idiosyncratic productivity. Firms are risk-neutral and consume the profits of production. When matched with a worker, firms use a homogeneous production technology, $y(\ell, \varepsilon)$, that converts labor hours, ℓ, into the economy's consumption good subject to an idiosyncratic productivity shock, ε. Let

$$y(\ell, \varepsilon) = e^{\varepsilon} \lambda \ell^{\alpha} \tag{1}$$

where the curvature parameter, $\alpha < 1$, captures the diminishing returns to scale from ignoring the fixed factors of production. The distribution of productivity shocks is firm-specific. Each period firm n receives an independent draw from $\varepsilon \sim N(\mu_n, \sigma_n^2)$ where I assume that $\mu_n = -\frac{1}{2}\sigma_n^2$ so that firms differ in the variance but not in the mean of their expected productivity shock.[5] Firms draw their type, σ_n, from density F_σ upon entry into the economy and their type stays fixed for the life of the firm. The firm's type determines the

[5]This fact follows from $\mathbb{E}[\exp(\varepsilon)] = \exp(\mu_n + \frac{1}{2}\sigma_n^2) = 1$. This assumption is stronger than the typical normalization when using a scale parameter since the scale parameter is not firm-specific. I discuss the (non-)identification of μ_n in Section IV.

variance in the marginal productivity of labor across periods and therefore their return from adjusting hours.

B. Workers

All workers have the same per-period endowment of hours which they divide between labor and leisure. In the production phase, employed workers receive compensation in return for hours worked. When unemployed, workers receive a common unemployment benefit b. Workers face a within period budget constraint such that consumption in any period must be less than or equal to compensation, $w\ell$, or b if unemployed. Equivalently,

$$C \leq w\ell + \mathbb{1}_{[ue]}b. \tag{2}$$

From the perspective of workers, contracts where hours vary across periods generate random sequences of consumption and labor hours and their inability to adjust their labor supply or borrow and save to smooth consumption and leisure represents the primary friction in the model. Let the utility function $Z(C, \ell)$ describe the preferences of a worker over consumption C and labor hours ℓ in any period as

$$Z(C, \ell) = \frac{C^{1-\rho_c}}{1-\rho_c} - \delta\frac{\ell^{1-\rho_\ell}}{1-\rho_\ell} - \mathbb{1}_{[ue]}g(z). \tag{3}$$

Utility is time-separable and exhibits constant relative risk aversion in consumption and leisure with δ as the relative weight given to hours. Workers are heterogeneous in their degree of risk aversion which is summarized by their type $\rho = \{\rho_c, \rho_\ell\} \in \mathbb{R}^2$ which has the bivariate distribution $\rho \sim F_\rho$ in the population. The $g(z)$ term captures an additional stigma cost of unemployment that depends on a vector of observable worker characteristics, z. I assume $g(z) > 0$ for all z which implies that all workers prefer employment, all else equal.

C. Optimal Labor Hours

In the production phase firms maximize profits subject to the agreed upon contract, namely a wage and whether they can adjust hours after seeing the period's production shock. A firm holding a variable contract with wage w_v solves

$$\max_{\ell} \; e^{\varepsilon}\lambda\ell^{\alpha} - w_v\ell, \tag{4}$$

implying a profit maximizing labor demand of

$$\ell_v(\varepsilon) = \left(\frac{\alpha\lambda e^{\varepsilon}}{w_v}\right)^{\frac{1}{1-\alpha}}. \tag{5}$$

In contrast, a firm in a fixed contract with wage w_f sets l_f before the ε realization and so maximizes expected profit

$$\max_{\ell} \ \lambda \ell^{\alpha} - w_f \ell \tag{6}$$

and therefore has a shock-independent labor demand of

$$\ell_f = \left(\frac{\alpha \lambda}{w_f} \right)^{\frac{1}{1-\alpha}}. \tag{7}$$

III. Bargaining

All newly met pairs negotiate contracts during the second phase in a full information environment such that each agent knows the other's type. The only uncertainty is over productivity shocks and whether the match will be exogenously destroyed. The firm proposes a wage and distribution of potential hours. The worker may choose to accept or reject this contract though neither party may terminate a signed contract in future periods. In the variable sector the firm's distribution of productivity shocks determines the distribution of potential hours. In the fixed sector hours are independent of the shock and the same across periods. In equilibrium, the wage and distribution of hours are set through firm optimization behavior described in the following section.

Either party may reject the match. Workers reject any match that does not satisfy their participation constraint. Firms reject a match if the expected (discounted) return from re-posting its vacancy next period is larger than the value of the match. Failed meetings result in firms receiving a pay-off of zero, workers receiving their unemployment benefit, and both parties entering the next period as unmatched. If neither party rejects the contract they enter the production phase and workers supply labor and receive compensation in accordance with their contract. All contracted hours and pay-outs are binding for both parties during the production phase.

I assume bargaining follows a version of the repeated infinite bargaining of Rubinstein (1982) when the worker's bargaining weight is set to zero. Cahuc et al. (2006) show this setting results in the expected profit maximizing firm making a take-it-or-leave it offer to the worker such that the value of working for a firm of type σ is exactly equal to the worker's reservation value. Formally, the contract must satisfy the worker's participation constraint

$$\mathbb{E}_{\varepsilon} Z(w_m \ell_m, \ell_m) + \beta[(1-s)J^m(\rho,\sigma) + sU(\rho)] \geq Z(b,0) - g(z) + \beta U(\rho),$$

where the firm's take-it-or-leave-it offer just satisfies the inequality and workers are left at their reservation value.

Given the bargaining environment, in equilibrium workers are *ex-ante* indifferent between working under a fixed contract, working under a variable contract, and remaining unemployed. In similar settings where workers are indifferent between working and not working, they are typically assumed to arbitrarily prefer working in order to provide a decisive policy function. Here, the worker faces three options and an analogous tie-breaking rule has serious implications for labor market sorting and, consequently, welfare analysis. To mitigate this issue I assume that firms are unable to contract on the stigma cost of unemployment, $g(z)$, and instead treat the worker's pay-off from unemployment as $Z(b,0)$. Equivalently, they act as though $g(z) = 0$ for all z. Therefore, the worker's effective participation constraint when bargaining becomes

$$\mathbb{E}_\varepsilon Z(w_m \ell_m, \ell_m) + \beta[(1-s)J^m(\rho,\sigma) + sU(\rho)] \geq Z(b,0) + \beta U(\rho). \tag{8}$$

This assumption provides an intrinsic premium of employment where firms leave money on the table even when z is fully known by both parties. A rationale for this assumption is that z contains worker characteristics like race, sex, national origin, religion, and disability, over which firms are legally prohibited from considering in labor contracts in the United States. These characteristics plausibly affect the value of unemployment but are assumed irrelevant for a worker's productivity and, were a firm to write contracts that extract this additional surplus, evidence of this behavior invites legal penalties for labor discrimination.

A. Optimal Labor Compensation

The optimal contract for a firm to offer during the negotiation phase leaves workers just indifferent between accepting the contract and their effective outside option. Workers know the firm's type and their rule for determining hours during the production phase. Therefore, the equilibrium wage offer during negotiation in sector m solves

$$\min_{w \in \mathbb{R}^+} w \quad \text{s.t.} \tag{9}$$
$$\begin{cases} \mathbb{E}_\varepsilon[Z(w\ell_v(w,\varepsilon), \ell_v(w,\varepsilon))] + \beta\Big[(1-s)J^v(\rho,\sigma) + sU(\rho)\Big] & \geq Z(b,0) + \beta U(\rho) \quad \text{for } m = v \\ Z(w\ell_f(w), \ell_f(w)) + \beta\Big[(1-s)J^f(\rho,\sigma) + sU(\rho)\Big] & \geq Z(b,0) + \beta U(\rho) \quad \text{for } m = f \end{cases}$$

where the relevant constraint enforces that the worker's expected value of working be at least as great as the worker's effective outside option. Given the strict concavity of utility the constraint binds at interior solutions. Re-arranging terms we have that the wage offer satisfies the generalized wage setting equation:

$$\mathbb{E}_\varepsilon[Z(w\ell(w,\varepsilon),\ell(w,\varepsilon))] = Z(b,0) + \beta(1-s)[U(\rho) - J^m(\rho,\sigma)]. \tag{10}$$

While expected profits are independent of firm type, curvature in the production makes expected utility a function of the firm's type. As an example, consider the expected hours under a variable contract:

$$\mathbb{E}_\varepsilon[\ell(\varepsilon)] = \mathbb{E}_\varepsilon\left[\left(\frac{\alpha\lambda\exp\{\varepsilon\}}{w_v}\right)^{\frac{1}{1-\alpha}}\right] = \left(\frac{\alpha\lambda}{w_v}\right)^{\frac{1}{1-\alpha}}\exp\left\{\frac{\alpha}{2(1-\alpha)^2}\sigma^2\right\}.$$

The non-linearity of (10) means a solution, much less a unique solution, may not exist for a given set of parameters. However, given a set of solutions the following lemma establishes the optimal wage for a given firm.

Lemma III.1. For any set of solutions to (10), the smallest value in the set is the expected profit maximizing wage and therefore the firm's optimal wage offer.

Proof. Firms choose w to maximize expected profit which we can re-write as

$$\begin{aligned}
\mathbb{E}_\varepsilon[\pi(\ell(\varepsilon),w)] &= \mathbb{E}_\varepsilon\left[e^\varepsilon\lambda\ell^\alpha - w\ell\right] \\
&= \mathbb{E}_\varepsilon\left[\ell\left(\frac{w}{\alpha}\left(\frac{\lambda\alpha e^\varepsilon}{w}\right)\ell^{\alpha-1} - wl\right)\right] \\
&= \mathbb{E}_\varepsilon\left[w\ell\left(\alpha^{-1}\ell^{1-\alpha}\ell^{\alpha-1} - 1\right)\right] \\
&= \mathbb{E}_\varepsilon\left[w\ell\left(\alpha^{-1} - 1\right)\right] \\
&= \mathbb{E}_\varepsilon\left[w^{\frac{-\alpha}{1-\alpha}}(\lambda\alpha e^\varepsilon)^{\frac{1}{1-\alpha}}\left(\alpha^{-1} - 1\right)\right]
\end{aligned}$$

where clearly expected profit is strictly decreasing in w for $\alpha \in (0,1)$. □

One interesting consequence of the above environment and the binding nature of the contracts is that outcomes for firms and workers are not always *ex-post* optimal given the realized shock. For example, with some probability firms operating with fixed labor contracts experience a productivity shock small enough to cause negative profit. Similarly, workers in variable contracts have no guarantee that the realized hours and compensation will provide more utility than the outside option of unemployment. Incorporating endogenous "quits" and "fires" into the model removes this potentially undesirable feature but distorts the primary welfare cost of interest. Fixed contracts inherit uncertainty absent from the variable sectors where hours freely adjust. Moreover, requiring satisfaction of participation constraints *ex-post* removes the entire welfare cost of varying hours by providing full insurance either by converting to a fixed hours contract or allowing the effective wage to vary with the realized shock.[6]

B. Search

A worker strictly preferring a fixed contract would only ever seek employment in the fixed sector. In this sense workers direct their search by choosing a sector which allows sorting based both on their own preferences and in response to differences in the equilibrium probability of finding a job.

After entry and exit, unmatched workers and firms meet through a frictional search process. Within a sector search is random. Search is costless for unemployed workers who may search in either sector but not both. When choosing between the sectors workers know the probability they will meet a firm in either sector as well as the distribution of firms participating in the sector. Each period firms with an unfilled vacancy must pay a cost at the beginning of the search phase to keep the vacancy open and must remain in the sector where they created their vacancy. Firms know the distribution of participating worker types in each sector and the probability of a contacting a worker. I assume that search in the two sectors functions independently but using the same matching technology. Let u_m and v_m be the measure of unemployed workers and vacancies, respectively, in sector m. An aggregate matching function, $M(u, v)$, governs the measure of meetings between vacancies and unemployed workers where a meeting is a chance for worker and firm to negotiate a labor contract. I assume $M(u, v)$ has a Cobb-Douglas form with constant returns to scale such that the measure of meetings is determined by

$$M(u, v) = \mu u^{\tau} v^{1-\tau}.$$

For notational convenience, let θ_m be the ratio of vacancies to unemployed workers in sector m, or a sector-specific notion of labor market tightness.[7] Accordingly, the probability an individual applicant meets a firm with a vacancy is $\phi_m = M(u, v)/u = \mu\theta^{1-\tau}$ while the probability that a particular firm with a vacancy contacts an unemployed worker is $q_m = M(u, v)/v = \mu\theta^{-\tau}$. Though determined in equilibrium, all agents take these probabilities as given when searching.

To formalize the search decision let $J^m(\rho, \sigma)$ represent a worker's value of entering the period as employed with firm-type σ in sector m and $U(\rho)$ be the value of entering the period unemployed. Workers have preferences over the realized hours and compensation during the production phase so the value of any match is the expected utility of the contract plus the expected value of continuing into the next period. For a

[6] For more on optimal contracting between risk averse workers and risk neutral firm see Cooper (1985).

[7] This convention captures the notion that in a tight labor market there are relatively fewer unemployed workers per each open vacancy which, for example, makes the creation of additional vacancies less attractive.

worker-type ρ in a variable contract with a firm-type σ this value is

$$J^v(\rho,\sigma) = \mathbb{E}_\varepsilon \left[Z(w_v \ell_v(\varepsilon), \ell_v(\varepsilon); \rho) \right] + \beta \left[(1-s) J^v(\rho,\sigma) + sU(\rho) \right] \tag{11}$$

where s is the probability the job is exogenously destroyed at the end of the period. Neither party may terminate an existing contract and all match destruction is exogenous. For workers holding fixed contracts all pay-offs are independent of the realized productivity shock as evidenced in

$$J^f(\rho,\sigma) = Z(w_f \ell_f, \ell_f; \rho) + \beta \left[(1-s) J^f(\rho,\sigma) + sU(\rho) \right]. \tag{12}$$

Consequently, the value of their match is also independent of the firm's type, *e.g.* $J^f(\rho,\sigma) = J^f(\rho)$, and their per-period utility pay-off is perfectly smooth for the duration of the match.

Let $M^m(\rho,\sigma)$ be the worker's value of a meeting a firm such that if both parties agree to the contract they get the value of the job, $J^m(\rho,\sigma)$, and the value of remaining unemployed otherwise. Equivalently,

$$M^m(\rho,\sigma) = \begin{cases} J^m(\rho,\sigma) & \text{if both parties accept} \\ Z(b,0) - g(z) + \beta U(\rho) & \text{otherwise} \end{cases} \tag{13}$$

Let $U^m(\rho)$ be the value of searching in sector m where with probability $\phi_m = \phi(\theta_m)$ the worker meets a firm so that

$$U^m(\rho) = \phi_m \mathbb{E}_\sigma \max\{M^m(\rho,\sigma), Z(b,0) - g(z) + \beta U^m(\rho)\} + (1-\phi_m)(Z(b,0) - g(z) + \beta U^m(\rho)). \tag{14}$$

The worker integrates the value of potential meetings over the distribution of unmatched firms. Having this notation, the value of unemployment for a worker of type ρ also describes the policy function for directed search in that the worker searches in the sector that satisfies

$$U(\rho) = \max\{U^f(\rho), U^v(\rho)\}. \tag{15}$$

At the beginning of a period, the set of vacancies in a sector is the set of firms entering the period unmatched plus new entries to the sector. Let $R(\rho,\sigma)$ be the value of a firm entering a period matched to a worker of type ρ. Again, pay-offs occur in the production phase so hiring decisions are made while integrating against the firm's own distribution of productivity shocks. Specifically, the value of a firm in the variable sector of type σ matched with a worker of type ρ is

$$R^v(\rho,\sigma) = \mathbb{E}_\varepsilon[y(\ell_v(\varepsilon),\varepsilon) - w_v\ell_v(\varepsilon)] + \beta\left[(1-s)R^v(\rho,\sigma) + sV^v(\sigma)\right]. \tag{16}$$

The value of being unmatched next period depends on the sector they initially entered as switching is disallowed. Analogously, the firm's value in the fixed sector is

$$R^f(\rho,\sigma) = \mathbb{E}_\varepsilon[y(\ell_f,\varepsilon)] - w_f\ell_f + \beta\left[(1-s)R^f(\rho,\sigma) + sV^f(\sigma)\right]. \tag{17}$$

I assume firms pay a sector specific cost of maintaining a vacancy, $\chi_m(\theta_m)$, that depends on the labor market tightness in the sector but not on the firm's own characteristics. Let $\chi_m(\theta_m)$ be strictly increasing in θ_m. Accordingly, the value of a firm of type σ entering the period unmatched in sector m is

$$V^m(\sigma) = -\chi_m(\theta_m) + q_m\mathbb{E}_\rho \max\{R^m(\rho,\sigma), \beta V^m(\sigma)\} + (1-q_m)\beta V^m(\sigma), \tag{18}$$

where firms integrate against the known distribution of unmatched workers in the sector, F_ρ^m, and take as given, $q_m = q(\theta_m)$, the probability of contacting a worker in sector m. Clearly, firms will reject any match where $\beta V^m(\sigma) > R^m(\rho,\sigma)$, or the value of continued search exceeds the value of the match. Intuitively, a high variance firm may pass on matching with a very risk averse worker and prefer to search again in the hopes of meeting a less risk averse worker.

C. Match Destruction

The set of unmatched agents entering the final phase are those who did not meet any potential partner plus those whose meetings failed to produce a viable contract. Subsequently, existing matches are destroyed with probability s which is independent of all firm and worker characteristics. Let $u_m(\rho)$ be the unemployment rate of worker-type ρ in sector m and $\eta_m(\rho)$ be the probability of meeting a compatible firm. The randomness of meetings within a sector implies that given the set of unmatched firms, the probabilities of contacting a firm and getting rejected are independent. The end of period unemployment for a worker of type ρ equals

$$u'_m(\rho) = s(1-u_m(\rho)) + (1-\phi_m\eta_m(\rho))u_m(\rho).$$

where $\eta_m(\rho) = \int_{F_\sigma^m} \mathbb{1}_{[R^m(\rho,\sigma)\geq\beta V^m(\sigma)]}$. Similarly, let $v_m(\sigma)$ be the measure of vacancies from firm-type σ in sector m such that

$$v'_m(\sigma) = s(1-v_m(\sigma)) + (1-q_m\psi_m(\sigma))v_m(\sigma),$$

where $\psi_m(\sigma) = \int_{F_\rho^m} \mathbb{1}_{[R^m(\rho,\sigma)\geq\beta V^m(\sigma)]}$, or the probability of a meeting resulting in a viable match.

D. Entry

At the beginning of any period any number of firms may enter either sector by paying a one-time cost of creating a vacancy in that sector. Concurrently, an existing firm may choose to exit the market and receive a pay-off of zero forever. Any worker matched with an exiting firm is treated as entering the period as unemployed. Firms learn their type after entry and may not enter or exit in any other phase. Once they have created a vacancy, firms cannot create a second nor switch to the other sector.[8] Formally, let Q be the value of entering and V^m be the value of holding a vacancy in sector m, then

$$Q(\sigma) = \max\{-c_E - c_v + V^v(\sigma), -c_E + V^f(\sigma)\} \tag{19}$$

where c_E is a common cost of entering and c_v is an additional cost specific to creating a vacancy in the variable sector. Since firms are free to exit we have the equilibrium condition that $V^m(\sigma) \geq 0$ for all firm types σ participating in sector m. An infinite number of potential firms may enter by paying the cost of creating a vacancy so in a stationary equilibrium we have the free entry condition $\mathbb{E}_\sigma[Q(\sigma)] \leq 0$ where the expectation is taken with respect to F_σ, the distribution of firm types. Taken together, these properties imply entry into each sector until the expected return from creating a position in either sector is zero. When deciding whether to enter, firms take market conditions as given.

E. Steady State Equilibrium

A steady state equilibrium requires the optimal behavior of agents taking as given both the exogenous components of the model and those determined in equilibrium as well as the consistency conditions enforcing stationarity.

The exogenous components of the model are the distributions of worker and firm types $\{F_\rho, F_\sigma\}$, the preferences parameters of the worker $\{F_\rho, \delta, g(\cdot)\}$, the labor elasticity and scale of production $\{\alpha, \lambda\}$, the scale and elasticity of the matching function $M(\cdot,\cdot)$ $\{\mu, \tau\}$, the entry costs $\{c_E, c_v\}$, the maintenance cost of a vacancy $\{\chi_f(\cdot), \chi_v(\cdot)\}$, the common unemployment benefit, b, the rate of job destruction, s, and the discount factor β.

The endogenous objects determined in equilibrium are (*i*) labor market tightness $\{\theta_v, \theta_f\}$, (*ii*) match acceptance rates for workers $\{\eta_f(\rho), \eta_v(\rho)\}$, (*iii*) match acceptance rates for firms $\{\psi_f(\sigma), \psi_v(\sigma)\}$, (*iv*) the

[8] Assuming firms pay a cost to learn their type and then choose which sector to enter and allows for firm sorting between the sectors. The identification of this sorting without demand side data relies heavily on the assumption that the expected productivity shock is independent of firm type.

distributions of unmatched agents in each sector $\{F^f_\rho, F^v_\rho, F^f_\sigma, F^v_\sigma\}$, and (*v*) the measure of participating firms in each sector determined in equilibrium by the free entry conditions.

The set of optimal firm and worker behavior outlined in (*i*) the entry decision of firms into the fixed and variable sector (19), (*ii*) the hiring decisions for firms (18), (*iii*) optimal labor contracts that maximize expected profit and satisfy participation constraints ((7), (5), and (9)), (*iv*) the worker's contract acceptance rule (8), and (*v*) the directed search choice of workers (15).

The assumption of stationarity implies that all inflows equal outflows for all worker and firms types whether measured by matched or unmatched status, type of contract, or sector participation. This property leads to following consistency conditions. Conditional on participation in sector m, the measure of unemployed workers of type ρ in sector m is

$$u_m(\rho) = \frac{s}{(s + \phi_m \eta_m(\rho))}. \tag{20}$$

Conditional on participation in sector m, the measure of vacancies from firm-type σ is

$$v_m(\sigma) = \frac{s}{(s + q_m \psi_m(\sigma))}. \tag{21}$$

In order to pin the dynamics of the economy in steady state equilibrium I first derive a few useful implied properties. Since the paper is motivated by applying counterfactual policy to an equilibrium with a mix of contract types I restrict analysis to the case where both sectors have participating workers and firms which motivates the following definition:

Definition III.1. Let sector m be *in operation* if the sector has a positive finite measure of unmatched firms and workers, or $\theta_m \in (0, \infty)$.

One consequence of the bargaining environment is that the *ex-ante* value of working is fully insured against the realized productivity shock, the firm's type, and even the contract type.[9] Substituting the participation constraint (8) into the value of working in sector m we have

$$\begin{aligned} J^m &= \mathbb{E}_\varepsilon Z(w_m \ell_m, \ell_m) + \beta[(1-s)J^m + sU] \\ &= Z(b, 0) + \beta(1-s)[U - J^m] + \beta(1-s)J^m + s\beta U \\ &= Z(b, 0) + \beta U. \end{aligned} \tag{22}$$

[9]When workers have positive bargaining weight they remain insured against the shock but pay-offs depend on the sector and firm type.

In other words, the equilibrium value of any job is equal to the pay-off from the common unemployment benefit plus the discounted value of being unemployed next period, both of which are independent of any characteristic of the current match aside from the worker's type. However, despite *ex-ante* equivalence, actual pay-offs may cause the *ex-post* values of jobs to differ substantially in any finite set of periods.

In the steady state equilibrium the stationarity of θ_m implies that for any given type ρ, if $U^f(\rho) > U^v(\rho)$ in one period then it will true for all periods and workers will always continue to search in the same sector next period. Consider the case where all firms are willing to match with workers, or $M^m = J^m$ and $\eta_m(\rho) = 1$ for all ρ. Workers direct their search to the sector satisfying $U = \max\{U^f, U^v\}$. If we substitute the value of a job (22) into the value of unemployment we get

$$
\begin{aligned}
U^m &= \phi_m \mathbb{E}_\sigma \max\{J^m, Z(b,0) - g(z) + \beta U^m\} + (1-\phi_m)(Z(b,0) - g(z) + \beta U^m) \\
U^m &= \phi_m \max\{Z(b,0) + \beta U^m, Z(b,0) - g(z) + \beta U^m\} + (1-\phi_m)(Z(b,0) - g(z) + \beta U^m) \\
U^m &= \phi_m \max\{Z(b,0), Z(b,0) - g(z)\} + (1-\phi_m)(Z(b,0) - g(z)) + (1 - \phi_m + \phi_m)\beta U^m \\
U^m &= \phi_m Z(b,0) + (1-\phi_m)(Z(b,0) - g(z)) + \beta U^m \\
U^m &= \frac{Z(b,0) - (1-\phi_m) g(z)}{1-\beta},
\end{aligned}
\tag{23}
$$

where the equilibrium value of entering a period unemployed in a particular sector is just the pay-off from unemployment less stigma cost of unemployment weight by the long-run expected portion of time spent in unemployment. For comparison, if firms could write contracts incorporating z then the value of unemployment is

$$U^m = \frac{Z(b,0) - g(z)}{1-\beta}.$$

Since entry is determined by $\max\{U^f, U^v\}$, (23) implies that $U^f = U^v$ if and only if $\phi_f = \phi_v$ where $U^f = U^v$ is a necessary condition to have both sectors operating. However, the assumption that firms never reject a match is restrictive as firms may prefer to take the pay-off of zero this period and re-post their vacancy next period. Intuitively, firms with greater variance in shocks in the variable sector may prefer less risk averse workers. The firm's value of a vacancy (18) embeds the firm's decision rules for accepting or rejects matches, *e.g.* $\max\{R^m, \beta V^m\}$.

Since search operates in a full information environment a worker of type ρ knows the probability they will meet a firm unwilling to match with them in sector m. Let this probability be $\eta_m(\rho) = \eta(\theta_m, F^m_\rho, \rho)$ which is a function of their own types, labor market tightness, and the distribution of other unmatched worker-types. Intuitively, the firm's expected return from continued search directly relates to the likelihood of another meeting $q_m(\theta_m)$ and the expected type to be encountered (F^m_ρ) which are both equilibrium objects.

Re-writing (23) to account for rejected matches we have

$$U^m(\rho) = \frac{Z(b,0) - (1 - \phi_m \eta_m(\rho))g(z)}{1-\beta}. \tag{24}$$

Since this relationship varies across ρ we no longer require $U^f(\rho) = U^v(\rho)$ for all ρ to have both sectors operating which in turn introduces scope for sorting between sectors.

Another useful property of the economy closely related to the equilibrium value of unemployment (24) is the worker's surplus from employment. A little algebra shows that the difference $J - U$ can be written as

$$\begin{aligned} J - U &= Z(b,0) + \beta U - U \\ &= (1 - \phi_m \eta_m(\rho))g(z) \end{aligned} \tag{25}$$

which is just the premium from employment weighted by the long run expectation of time spent unemployed. The surplus disappears if finding a viable job is a certainty or if $g(z) = 0$. The negative of this term shows up prominently in the wage setting equation (10) where intuitively the wage is increasing in the size of the stigma cost, though at a rate decreasing in the probability of finding a viable match. In this manner, the firm does recover some portion of the "excess" worker surplus even though in any given period the contracted pay-off is independent of z. As $g(z)$ increases so does the intrinsic and protected premium employment provides to the worker. As labor market tightness θ_m increases, the job contact rate $\phi_m(\theta_m)$ goes up which also raises the value of unemployment, $U^m(\rho)$ and by the wage setting equation (10) and (25) increases the required wage offer.

The stability of search dynamics depends on the optimizing behavior of workers and firms which includes sorting by workers, the entry (and exit) decision of firms, and the endogenous distribution of unmatched agents. In general, a worker may be rejected by all firms in the economy which creates an ambiguity over whether such a worker is *unemployed* or *discouraged*. This distinction has bite as the set of unmatched workers looking for employment is a key driver of dynamics in an equilibrium search model and so there may be multiple equilibria with different sets of discouraged workers. In the fixed sector, a firm's value of a match and vacancy are independent of the firm's type and so a worker rejected by one firm will be rejected by all firms. To avoid this uncertainty, I assume that for any worker in the fixed sector, at least one firm type, and therefore all, are willing to match which prohibits the occurrence of an unemployable worker. This condition amounts to enforcing a positive match surplus in estimation and is not without loss of generality when taking the model to data as the distribution of worker heterogeneity matters when fitting the observed profile of hours and wages.

Given the restriction on match surplus, firm entry is set by the free entry condition where given a distribution of unmatched workers in the steady state, as firms enter the labor market tightness θ_m necessarily increases driving up the vacancy contact rate $\phi_m(\theta_m)$ and down the worker contact rate $q_m(\theta_m)$. Substituting the value of holding a vacancy (18) into the free entry condition and re-arranging yields

$$q_m(\theta)\mathbb{E}_\sigma\mathbb{E}_\rho \max\{R^m(\rho,\sigma), \beta V^m(\sigma)\} + (1 - q_m(\theta))\mathbb{E}_\sigma \beta V^m(\sigma) \leq c_m + \chi_m(\theta) \tag{26}$$

for $m = v, f$. This condition limits the value of entry so that in equilibrium an operating sector contains only a finite set of firms.

Lemma III.2. Suppose $\gamma \in (0, 1)$ and market m is operating. Given a steady state distribution of unmatched workers, F^m_ρ, the solution to (26) is unique and the constraint binds.

Proof. That the constraint binds follows from the steady state assumption and that no firm would enter when the value of entry was less than the expected cost $c_m + \chi_m(\theta)$ since creating a position has less expected value than the outside option of zero forever. That the solution is unique follows from the monotonicity of the first derivative where all terms are functions of θ

$$\frac{\partial q_m}{\partial \theta}\mathbb{E}_\sigma\mathbb{E}_\rho \max\{R^m, \beta Q^m\} + q_m \frac{\partial \mathbb{E}_\sigma\mathbb{E}_\rho \max\{R^m, \beta Q^m\}}{\partial \theta} - \frac{\partial q_m}{\partial \theta}\mathbb{E}_\sigma \beta Q^m + (1 - q_m \frac{\partial \mathbb{E}_\sigma \beta Q^m}{\partial \theta} \leq \frac{\partial \chi_m}{\partial \theta} \tag{27}$$

If the sector is operating it must be the case that for some σ $\mathbb{E}_\rho R^m(\rho,\sigma) \geq \beta Q^m(\sigma)$ or firms would exit in the first phase and $\theta_m = 0$. Therefore, the LHS is a convex combination between a larger and a smaller value with weight q_m. Since this weight is decreasing in θ_m, as firms enter the LHS gets smaller, further compounded by the fact that required wage goes and monotonically decreases expected profits as in Lemma III.1, and consequently $R^m(\rho,\sigma)$ $\forall \sigma$. Clearly, as $\theta_m \to \infty$, $q_m \to 0$ and unmatched firms exit for $\chi_m > 0$. Further, $\frac{\partial \chi_m}{\partial \theta} > 0$ which serves to close the difference more quickly. That the solution is unique follows from the monotonicity of $q_m(\theta)$ for $\gamma \in (0, 1)$ and $\frac{\partial R^m(\rho,\sigma)}{\partial \theta}$. □

IV. Identification

Solving for the model's equilibrium requires specifying the set of exogenous components relating to worker type and preferences, firm type and technology, and search dynamics. These primitives map to the structural parameters summarized in Table 1 and provide a way to disentangle the consequences of policy for welfare and efficiency. The wage setting equation (10) involves parameters from each category and illustrates the challenges in separately identifying the effect of policy using only data with observed wages and hours. The

Table 1: Model Primitives

Primitive	Description	Estimation Method
Worker Type and Preferences: $\frac{C^{1-\rho_c}}{1-\rho_c} - \delta\frac{\ell^{1-\rho_\ell}}{1-\rho_\ell} - \mathbb{1}_{[ue]}g(z),\ (\rho_c, \rho_\ell) \sim (m\rho, s\rho)$		
δ	Utility weight on labor hours	Vignettes
$g(\cdot)$	Stigma cost of unemployment	Vignettes
$m\rho$	Mean vector for worker types	Vignettes
$s\rho$	Covariance matrix for worker types	Vignettes
Firm Type and Technology: $y = e^{\varepsilon}\lambda\ell^{\alpha},\ \sigma_n \sim (m\sigma, s\sigma)$		
α	Labor output elasticity	SMM
λ	Scale of production	SMM
$m\sigma$	Mean of firm types	SMM
$s\sigma$	Variance of firm types	SMM
Search and Matching: $M = \mu u^{\tau} v^{1-\tau}$		
τ	Elasticity of M w.r.t. u	SMM
μ	Scale of M	SMM
c_E	Common entry cost	SMM
c_v	Additional entry cost for variable market	SMM
$\chi_f(\cdot)$	Vacancy cost for fixed market	SMM
$\chi_v(\cdot)$	Vacancy cost for variable market	SMM
b	Unemployment benefit	SMM
Other		
s	Probability of job destruction	Calibrated
β	Discount factor	Calibrated

Table 2: A list and description of model primitives and the approach to estimation. Vignettes refers the first stage estimation using stated preferences from a discrete choice experiment. SMM stands for the second stage simulated method of moments estimator that takes first stage estimates as given when simulating the model equilibrium and matching observed moments from equilibrium jobs and worker flows.

standard approach adds firm or matched employee-employer panel data to incorporate additional variation as in, for example, Postel-Vinay and Robin (2002), Flinn (2006), and Lise et al. (2016), though in each case preferences are limited to the fixed wage a job offers. Cooper et al. (2007) develop a search model where risk neutral multi-worker firms fully insure risk averse but homogeneous workers against uncertain hours and compensation and use panel data on firms to identify firm technology and heterogeneity for a set of calibrated preference parameters. In Dey and Flinn (2005) a contract is a wage and whether the employer provides health insurance and workers vary in how much they care about the latter. They use the incidence of observed jobs with employer-provided health insurance and paired wages to identify the heterogeneity in preferences but model the firm's cost of provision as an exogenous random variable. In contrast, given the premise in my model that observed hours are not necessarily optimal hours from the perspective of the employee the actual sequence of hours provides limited information about worker preferences. However, given wages, they are informative about the firm's production technology. The primary challenge becomes distinguishing how preferences, firm technology, and aggregate market conditions precisely enter into wage and hours determination.

I propose an identification strategy novel to the equilibrium search literature that permits not only multi-attribute jobs but also worker heterogeneity in endogenous wage determination. The approach first isolates the preferences of workers in a controlled non-market environment and uses those estimated preferences as inputs in a second stage of estimation that involves solving for equilibrium labor contracts. Specifically I use an experiment where participants see a sequence of choice sets containing hypothetical jobs that map directly to the multi-attribute contracts of my model. The use of stated preferences has been used in the search literature by Bloemen (2008) who uses "desired hours" to estimate a model where employees are not allowed to choose their hours. van der Klaauw (2012) incorporates the expected future occupation of respondents when estimating a dynamic discrete choice model and shows the addition improves the efficiency of the estimator. I view my approach as one solution to the difficulties of identifying equilibrium search models without firm data, as described in Eckstein and Wolpin (1990) and Flinn (2006), and as a strategy particularly well-suited to models where jobs have multiple attributes and their specific value to firm and worker are of primary interest. The method requires an assumption about how these attributes enter worker preferences and firm profit which is largely a scenario-specific issue. Given preferences, the variation in the hours of observed jobs and worker flow moments pin down firm technology and search dynamics much the same way as in the equilibrium search papers cited above.

The following sections consider the three categories of structural parameters and further develop my identification strategy while the final section discusses the challenge in identifying the bargaining weight of workers with data from only one side of the market.

A. Preferences

To identify the primitives underpinning heterogeneous preferences in the model I use the Understanding Work Schedule Forecastability (UWSF) survey. The UWSF is a survey instrument designed by myself and Flávio Cunha to estimate the willingness to pay for the kind of job attributes in the Work Schedule questions from the NLSY97. The crucial section contains a series of vignettes where on each vignette respondents choose between three distinct options—two job offers and an unemployment benefit. Each respondent sees the same first vignette with jobs set to resemble entry-level job offers and thirteen subsequent vignettes randomly selected from a pool of fifty-one. The randomization balances the desirability of getting multiple responses to the same vignette with the need to sufficiently vary offered contracts and capture the correlation between attribute values. Figure 3 presents an example vignette from the survey. In the first column the job has no variation in weekly hours and an implied wage rate of $21.00. The second job has hours that vary between 20 to 60 hours per week though the wage rate here happens to be the same. Respondents weigh those jobs relative to a $1,680 unemployment benefit with no additional conditions. By varying the outside option the survey captures the decision to participate in the work force while choices between jobs reveal preferences over consumption and leisure. The remainder of this section describes the construction of the survey instrument. Information about the sample used for estimation is provided in Section VI.

The use of vignettes to trace out preferences has a long history in the marketing literature but has recently emerged in labor economics to measure the willingness to pay for job attributes. Hanley et al. (1998) shows this type of discrete choice experiment performs better than traditional stated preference methods when estimating willingness to pay. Because vignettes allow the econometrician to see the full set of attributes for each option, the set of observed choices map out the preferences for any given respondent. Eriksson and Kristensen (2014) use the choices of Danish employees to estimate willingness to pay for fringe benefits like flexibility in scheduling, training, and health insurance and even suggest that human resource departments may find this method useful in reducing informational frictions in contract negotiations. Mas and Pallais (2016) embed such an experiment in a recruiting process and non-parametrically estimate the willingness to pay for attributes like scheduling flexibility and the opportunity to work from home by exogenously varying the wage offer.

To make the hypothetical jobs as realistic as possible, the initial stage of the survey prompts respondents to suppose they receive twenty job offers over the next three months and provide the the largest, smallest, and mostly likely wage offers to occur. The vignettes contain four possible wage values created from convex combination of these person-specific moments. Even if the respondents vary the non-wage job attributes in their thought experiment they are still tracing out a distribution of likely wage rates that allow the vignettes

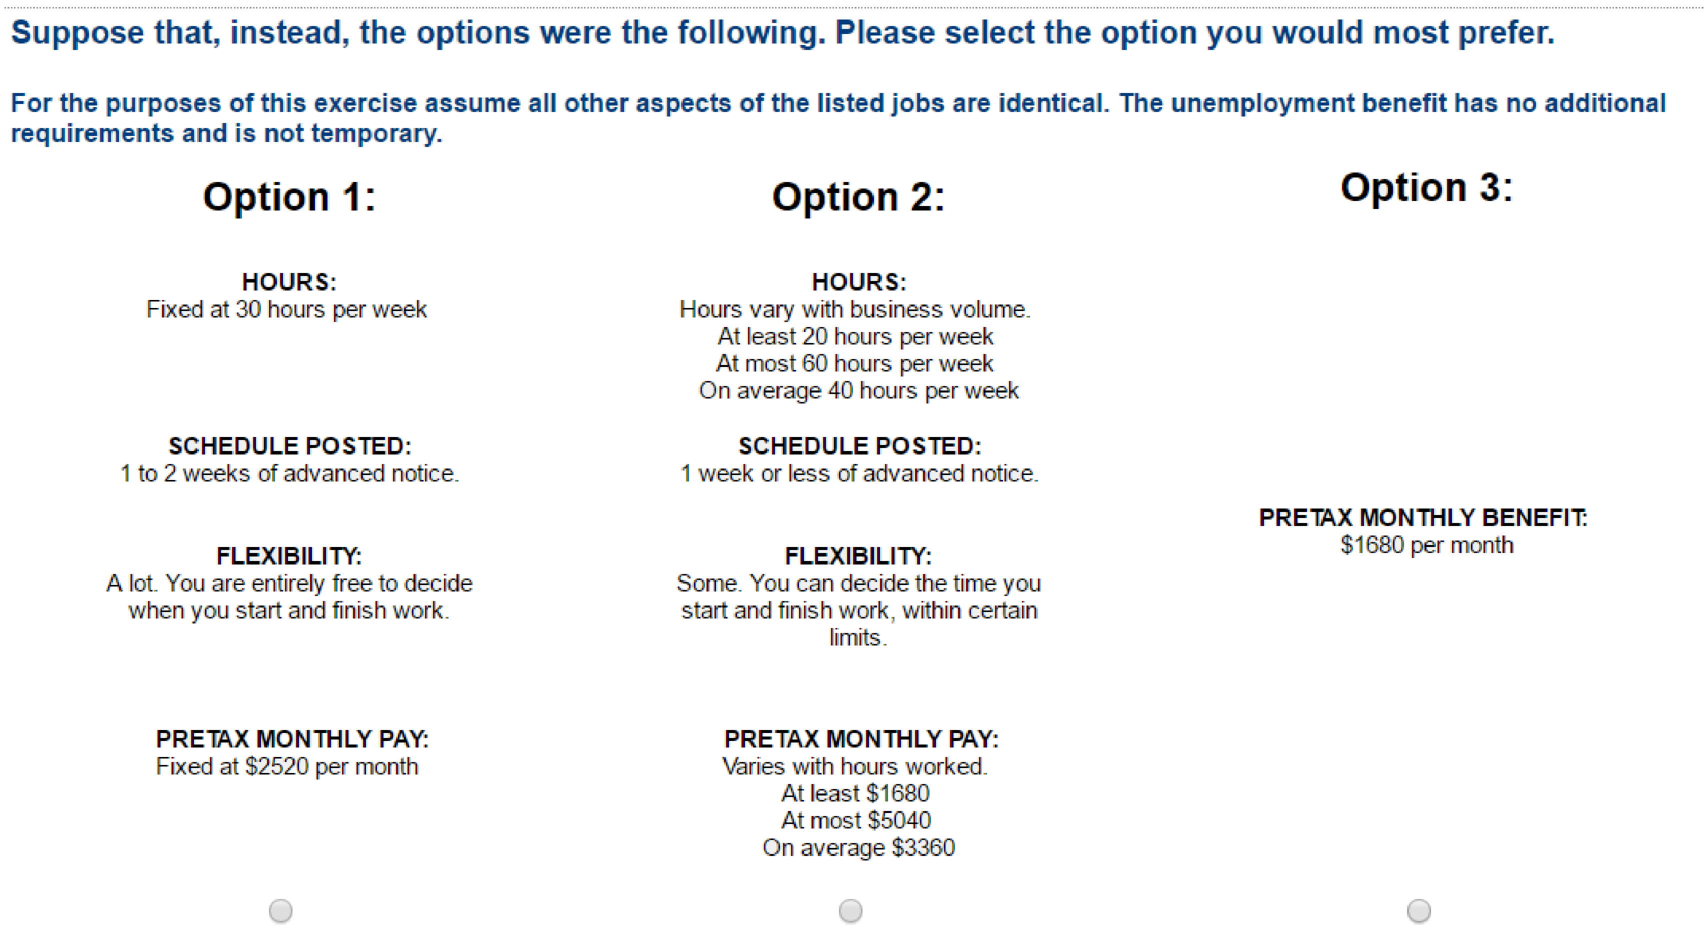

Figure 3: An example vignette from the UWSF survey instrument used to estimate worker preferences. Each respondent saw one vignette with entry level type jobs and an additional 13 randomly selected vignettes to trace out their preferences across the set of attributes as well as their participation decision. The wage offers are tailored to the individual though all other attributes are specific to the vignette. Appendix B provides additional details.

to tailor the wage offers and unemployment benefits to the individual. Notably, this feature also permits unemployed respondents to participate and thus potentially captures work force compositional effects within the larger model. Appendix B contains details of how the hypothetical wages are constructed from the expected wage offers.

All jobs are fully characterized by six dimensional attribute vector that describes hours, variation in hours, wage rate, variation in wage rate, flexibility, and degree of advanced notice. Each dimension has a discrete set of possible values. For example, a job may provide the worker with less than one week, two to three weeks, or three or more weeks of advanced notice of their schedule. For the purposes of the survey a job is vector with one value from each of the six possible dimensions where Table 3 describes the set of possible values. By constructing a closed set of possible attributes the econometrician generates a variety of choice sets that exogenously vary compensation while observing the entire set of alternative-specific characteristics and the outside option. This approach also permits assessment of combinations that may not exist in the equilibrium market, particularly for a given wage. The inclusion of such jobs is important as the absence of such a job does not provide any information on whether the job is dis-preferred by workers, firms, or both. In this way, using vignettes overcomes many of the standard issues when recovering willingness to pay from observed jobs like unobserved outside options, unobserved attributes, and potentially insufficient variation.

The vignettes vary both the effective wage rate and the distribution of possible hours independently which provides identifying information about the risk aversion over consumption, ρ_c, and over labor hours, ρ_ℓ, while the relative weight δ is pinned down by variation in levels. The outside option on each vignette, described as an unemployment benefit with "no additional requirements" and as "not temporary," represents a version of the model's common unemployment benefit though nothing in the survey disciplines the offered benefit to be an equilibrium value. Therefore, I do not use the vignettes to estimate b but respondents who choose an obstensibly worse job over the benefit reveal the extent and magnitude of the stigma cost of unemployment, $g(\cdot)$. The apparent disinclination is modeled as a utility stigma cost of unemployment that may vary across individuals. The degree to which the stigma cost influences decisions is an empirical question. Finally, since the full set of attributes is observed, variation in observed choices within a random utility framework can be decomposed into specific sources. Variation in observed choices within a vignette traces out preferences across job attributes, variation across respondents pins down the degree of heterogeneity in such preferences, and anything else is a person, vignette, and alternative specific preference shock.

Table 3: Job Attributes in the UWSF Vignettes

Dimension	Values	Notes
Hours	10	hours per week
	20	hours per week
	30	hours per week
	40	hours per week
Variance in Hours	Fixed	e.g. always $\overline{h}$ hours/week
	$\pm 20\%$	Hours vary between $.8\overline{h}, \overline{h}, 1.2\overline{h}$
	$\pm 50\%$	Hours vary between $.5\overline{h}, \overline{h}, 1.5\overline{h}$
Pre-Tax Hourly Wage	Low	$0.9 \times w_{low} + 0.1 \times w_{mid}$
	Mid1	$0.6 \times w_{low} + 0.4 \times w_{mid}$
	Mid2	$0.4 \times w_{mid} + 0.6 \times w_{hi}$
	High	$0.1 \times w_{mid} + 0.9 \times w_{hi}$
Variance in Pay	Fixed	Always w per hour.
	$\pm 10\%$	The effective wage rate varies between $.9\overline{w}, \overline{w}, 1.1\overline{w}$
Advanced Notice	$\leq$ 1 week	"1 week or less of advanced notice"
	1–2 Weeks	"1 to 2 weeks of advanced notice"
	3+ weeks	"3 weeks or more of advanced notice"
Scheduling Flexibility	None.	"Starting and finishing times are decided by your employer and you cannot change them on your own"
	A little.	"Starting and finishing times are decided by your employer but with your input"
	Some.	"You can decide the time you start and finish work, within certain limits"
	A lot	"You are entirely free to decide when you start and finish work"

A breakdown of the possible attributes for a job within a vignette. A job is characterized by having one value from each category. The description of the job is then an implied description of those attributes. For example, a distribution of compensation since wage and/or hours may vary. For a description of how wages were set see Appendix B.

B. Labor Contracts

Given worker preferences and heterogeneity, I use a set of equilibrium jobs from the NLSY97 described in Section V to identify the firm's primitives. Respondents report a wage and distribution of hours that I interpret as a wage and a sample of realized hours originating from a labor contracting process equivalent to that in the model. In other words, workers accept a job knowing the distribution of potential hours and the wage which are then fixed for the duration of the contract. Intuitively, since the wage is fixed before the hours are set, the variation in hours within a contract relates directly to the shocks that firm faced. This variation provides information about that specific firm's type as can be seen by looking at the equation for variance in log hours given by

$$\text{Var}(\ln \ell) = \frac{\sigma_n^2}{(1-\alpha)^2}.$$

Variation across contracts provides information about both the distribution of firm types as well as the common production technology. In equilibrium all firm types participate in the variable market so F_σ with parameters $m\sigma$ and $s\sigma$ is directly related to the variance of hours in observed jobs. However, the set of observed contracts is not random but rather a product of equilibrium sorting between workers and firms. Even if we could condition on the set of worker types that participate in the variable market, not all firm-worker pairs form viable matches which further motivates the solving of the general equilibrium model which enforces the selection effect in estimation.

To a first order approximation many of the production parameters relate directly to the observed levels in wages and hours. The unemployment benefit b, the scale of the production function λ, and the elasticity of labor α, contribute to the wage and hours level for any given firm-worker pair. Firm technology, in α and λ, further determines the location of the distribution of hours for a given wage. These relationships are non-linear but well-defined and so parameter values can be inferred through a comparison between the jobs in the simulated equilibrium and observed jobs.

One limitation of estimating the model without firm level data is the restriction that $\mu_n = -\frac{1}{2}\sigma_n^2$. The identification of a separate firm type without strong parametric assumptions is tenuous. Consider the expected hours of a variable contract which with some manipulation can be written as

$$\mathbb{E}_\varepsilon \left[\ell_v(w, \varepsilon)\right] = (\alpha\lambda)^{\frac{1}{1-\alpha}} w^{\frac{-1}{1-\alpha}} e^{\frac{\mu_n}{1-\alpha}} e^{\frac{\sigma_n^2}{2(1-\alpha)^2}}.$$

The mean of the firm type in μ_n is not separately identified from λ and can be safely normalized to unity but this normalization does not provide identification of the distribution of μ_n. Since we cannot condition equilibrium jobs on the type of worker, without some fixed relationship between σ_n and μ_n multiple combinations of w and μ_n generate the same level of expected hours. The standard econometric solution to this issue uses data where one firm has multiple workers which provides a way to recover the firm-specific productivity. Lacking this data or a strong prior to pin down this relationship I make the simplification that $\mu_n = -\frac{1}{2}\sigma_n^2$.

C. Search

The third ingredient necessary to map equilibrium jobs to the optimal labor contracts of the model is the relationship between search frictions and wage determination. The value of employment over remaining unemployed appears in the wage setting equation (10) which in the steady state equilibrium simplifies to a function of vacancy contact match acceptance rates as in (24). The contact and firm acceptance rates are governed by an aggregate matching technology that though treated here as structural parameters are perhaps better thought of as a reduced form approximation of market dynamics. For any given set of firm technology parameters, the matching technology and search costs largely determine the level of firm entry into a sector which in turn drives worker sorting through the choice between labor market sectors. My approach to estimating these parameters resembles that in Cooper et al. (2007) and Lise et al. (2016) who discipline the equilibrium to match observed moments from hiring, separations, vacancy creation, and unemployment in the labor market.

While differences in the matching probabilities of the two sectors drive sorting between sectors, within each sector the wage setting equations involve the common probability of exogenous match destruction, the sector-specific vacancy contact rate, and the sector- and worker-specific match acceptance rate. The probability of match destruction is independent of all firm, worker, and match characteristics and in a steady state equilibrium must exactly equal the flow of new matches per period. I directly observe the latter in the data as the measure of monthly hires in the retail sector in the Job Openings and Labor Turnover Survey. I treat this rate an a known quantity and calibrate the probability of job destruction outside of estimation. A analogous calculation of the vacancy contact rate is thwarted by the possibility that firms reject some matches and failed meetings are not recorded. Moreover, the relationships between the aggregate moments in the data and sector-specific moments of the labor market model complicate a straightforward calibration. For example, the available data does not contain separate records of vacancies where hours vary and where they do not. Therefore, the observed rate of vacancies corresponds to the rate of simulated vacancies aggregated across the two sectors. Identification relies on how search and search frictions influence unemployment levels, the level of vacancies, unexplained differences in labor contracts, and indirect inference on worker sorting. For example, the scale of the matching function largely sets the level of unmatched firms and workers in each sector which in turn imply aggregated rates of steady state unemployment and vacancies. The elasticity of matches with respect to unemployment plays a role in sorting dynamics as the more elastic the technology the greater the scope for large movement of workers between the markets for a comparatively small change in labor market tightness.

In contrast, while the matching technology pins down dynamics for a given labor market tightness, the costs of search set the level of equilibrium labor market tightness by controlling firm entry. Firms decide whether to enter knowing the distribution of unmatched workers in a sector and take contact rates as given, to the cost of entry, c_E, determines the measure of entering firms but also, indirectly, the distribution of participating workers. This property provides a mechanism to drive worker sorting and the observed ratio of variable to fixed contracts. Conversely, the costs of maintaining a vacancy provide a way to fine-tune the sector specific labor market tightness as the costs are a function of labor market tightness. The differential between the vacancy contact rates is a primary driver of worker sorting and so, given preferences, is partially disciplined by unexplained differences between labor contracts in the two sectors.

An important distinction in relating the primitives underlying search dynamics in the model to the moments from the retail sector is the implication that observed moments are in fact aggregations across the two sectors. For example the notion of aggregate labor market tightness is in fact $\theta = \frac{u_v+u_f}{v_v+v_f}$ which does not immediately map to the equilibrium objects θ_v and θ_f. While focusing on the retail industry increases the plausibility of the model, one drawback of considering only one sector is the lack of industry-specific unemployment rates. Unemployment not only helps pin down contact rates in the model but also the notion of labor market tightness. The labor market tightness proves an important moment to identifying the scale and elasticity of the matching function $M(u, v)$. Lacking a less objectionable alternative, I assume the retail market has the same unemployment rate, and consequently aggregate labor market tightness, as the larger U.S. labor market.

D. A Note on Worker's Bargaining Power

Setting the worker's bargaining power to zero is not without loss of generality. Assuming workers have a positive bargaining weight provides them with a positive share of match surplus through the wage setting equation. The primary complication is that in an equilibrium setting it produces a premium to employment that varies with firm and contract type. This feature removes the necessity of the exogenous premium provided by $g(z)$ but significantly complicates search dynamics and policy functions. Without bargaining weight, the (expected) value of a job is independent of productivity shocks, firm type, and even contract type which means the distribution of unmatched firms in a particular market is effectively irrelevant to the worker's decision. When workers receive part of the match's surplus this distribution becomes important to the directed search decision and solving for the equilibrium requires simultaneously solving for the distributions of unmatched firm and worker types. The directed search decision of workers and the consequent match rejection rules for both firms and workers create endogeneity between these objects and may imply multiple equilibria. Setting the parameter to zero simplifies the solution and provides a unique equilibrium.

A related concern is the identification of the parameter given the data used here. In principle, the structural parameters of firm type and technology as well as the bargaining weight would be identified by variation in the wage and hour profiles of equilibrium jobs. Since all three sets of primitives affect the level of wages and hours, conceivably for any value of the worker's bargaining parameter a set of parameters, particularly $\{b, \lambda, \alpha, m\sigma\}$, could be found to rationalize observed levels. Identification relies on the variation of surpluses across matches to isolate the bargaining weight both from a common outside option and the heterogeneity in the firm's type. The literature has long recognized the difficulty of identifying the parameter using only supply-side data (see Eckstein and Wolpin (1995) and Flinn (2006)). Flinn (2006) and Dey and Flinn (2005) estimate models without labor hours and use labor's share of total revenue from demand side data as an informative moment about the bargaining weight. In my model labor's share equals the output elasticity of labor α by virtue of the Cobb-Douglas specification of production which limits the usefulness of this approach. Given the high rate of turnover in the retail sector, a different tactic is to add on-the-job search which the literature suggests has more explanatory power when trying to match the wage profile within and between employment spells (see Lise et al. (2016)). This addition would increase the outside option of workers but adds considerable complexity to the steady state distributions of worker flows and matches and requires additional assumptions to extend to an environment with two separate sectors.

To provide a reference value on the possible range of the parameter, perhaps the most relevant estimates come from Flinn (2006) and Cahuc et al. (2006). Flinn looks at the U.S. labor market population that would be affected by the minimum wage and finds a precisely estimated bargaining weight of 0.42. This estimate may represent an upper bound if his model, by ignoring the non-wage job attributes of minimum wage jobs used for estimation allows the bargaining weight to capture compensating differentials that differ across jobs and so would not be fully accounted for by a common unemployment benefit. Cahuc et al. estimate bargaining weights for four sectors and find the uniformly lowest values for the service industry where they are approximately zero. However, their model incorporates on-the-job search and estimates parameters using French data which limits the direct comparability. Evidence presented in the next section shows the observed retail sector jobs typically have limited control over scheduling and relatively low wage rates which may suggest a lack of bargaining power.

V. A Sample of Retail Sector Employees

The primary sample in this paper consists of retail workers from Round 16 of the National Longitudinal Survey of Youth 1997 Cohort (NLSY97) which was conducted between November 2013 and July 2014. Respondents were born between 1980 and 1985 and range in age from 28 to 33 years at the end of 2013. I filter the sample to respondents who are not self-employed and report their current or most recent job as being in the retail sector with a 2002 Census Industry Codes between 4670 and 5580.[10] The bulk of these jobs are with department stores, supermarkets, and specialty stores. Table 10 contains a more detailed breakdown of the retailers by census code. Filtering for a complete set of job attributes and demographic information leaves 405 jobs characterized by a wage and a distribution of hours. I assume these jobs represent observations from a single labor market equilibrium with dynamics and structure equivalent to that of the model.

I view the retail sector as maximizing the plausibility of my model of labor contracting while minimizing additional un-modelled features such as overtime, unions, and productivity pay. I ignore minimum wage considerations as the common unemployment benefit should capture some degree of bunching at low wage values and also to avoid a zero likelihood issue as around 9% of observations have an imputed wage smaller than the 2013 federal minimum wage of $7.25. To avoid incorporating observations where the wage was contractually random I use the provided measure of hourly wage excluding tips, commissions, and overtime pay. For seventy-nine percent of the sample there was no difference between this measure and the one containing additional compensation.

The NLSY97 has always asked respondents their usual hours per week but beginning in Round 15 respondents were further asked to report what were the most and fewest hours they worked over the last four weeks. These three moments typify a device commonly used in surveys to succinctly capture the main features of a distribution. They are presumably easy for respondents to identify and readily translate to a Triangular distribution. I use these moments to construct my measure of week to week variation in hours. Accordingly, any job with all three moments the same has a fixed schedule. Since hours in the model are a continuous variable with support over the positive real line I adapt this reported measure into a continuous distribution with the same support. I assume usual, fewest, and most hours are sample statistics from some finite set of realizations for labor contracts generated in accordance with the model. I then make use of the convenient fact that independent and identically distributed realizations from a Triangular distributions are well-approximated by a Normal distribution with the same mean and variance as the Triangular distribution. The degree of approximation is illustrated in Figure 4.

[10] I exclude observations with census industry codes between 5580–5790, which includes electronic retail, vending machine operators, and direct selling to remove jobs I view as likely to have a qualitatively different type of labor contract.

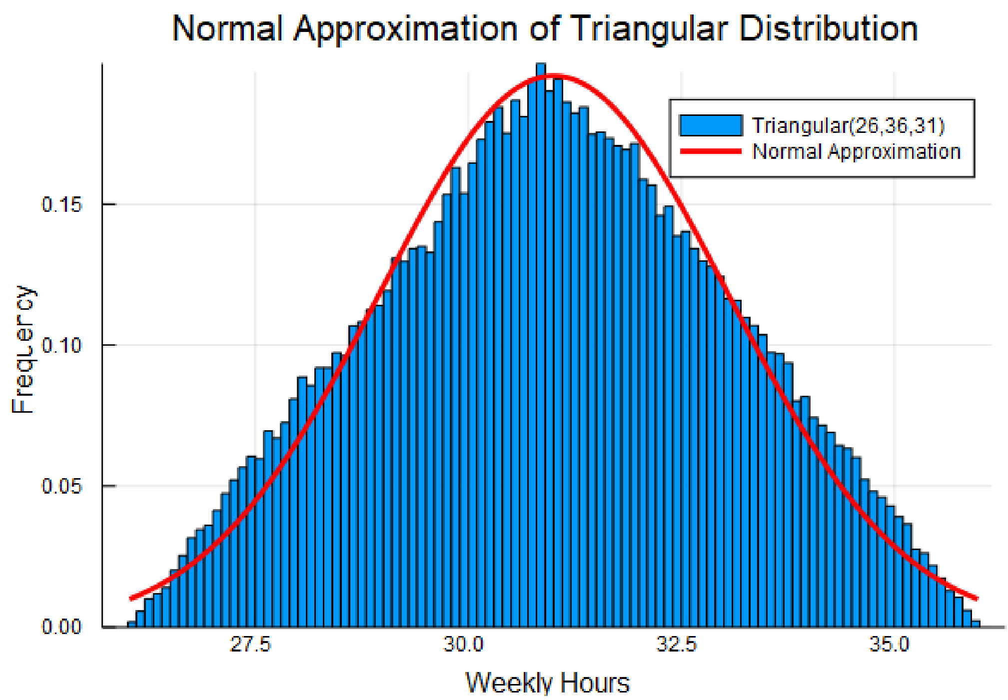

Figure 4: The weekly hours described by a Triangular distribution are well-approximated by a Normal distribution with the same mean and variance. The histogram has identically and independently distributed realizations from a Triangular distribution where hours range between 26 and 36 hours per week with 31 being the most likely.

For tractability my model uses the friction created by employer discretion over weekly hours as the source of unpredictability in labor scheduling. I motivate this approach by looking at jobs in the retail sector where workers have little control over hours. Other labor markets that exhibit as much if not more variability in hours would require a decidedly different model (see Barlevy and Neal (2017)). Moreover, the literature identifies many possible sources of instability other than variation in hours. Evidence from surveys and discrete choice experiments discussed earlier suggest that the degree of flexibility in the timing of hours and advanced notice of schedule play an important role in the valuation of a particular job. Intuitively, both of these attributes mitigate the scheduling burden of not knowing the number of hours they will work in any particular week but do not innately aid in smoothing consumption and leisure across periods. For this sample, eighty-two percent of employees have at most some input but no direct control over when they work. Similarly, when you condition on employees with schedules that actually vary from week to week then sixty-eight percent have at most two weeks of advanced notice of their schedule. Finally, sixty percent of the sample have together at most two weeks advanced notice and at most some input into their hours schedule. I interpret this evidence as consistent with the labor contracting framework in the model where workers have little bargaining weight or control over the exact number of hours in any week.

Table 4: Characteristics of Observed Retail Jobs

	Variable			Fixed			Total (Weighted)		
	Mean	Median	Std. Err.	Mean	Median	Std. Err.	Mean	Median	Std. Err.
Wage									
Wage	14.81	12.0	9.0	14.0	11.08	9.85	14.83	12.0	9.55
No Reported Tips, etc.	77%			81%			79%		
Tips, etc. as % of Wage	31%	7%	0.58	13%	10%	0.13	23%	8%	0.46
Hours									
Usual Hours	36.63	40.0	10.45	35.97	40.0	10.7	36.38	40.0	10.73
Hours vary	100%			0%			55%		
Hours $<$ 35	32%			24%			28%		
Hours $>$ 40	19%			6%			15%		
Most Hours in a Week	42.42	40.0	11.39				39.6	40.0	11.74
Fewest Hours in a Week	31.65	35.0	11.71				33.7	39.0	11.3
Instability Ratio (IR)	0.34	0.25	0.32				0.19	0.08	0.3
IR (Hours < 35)	0.58	0.4	0.44				0.36	0.2	0.48
Overtime Hours $>$ 0	12%			4%			8%		
Overtime Hours (if $>$ 0)	4.0	2.5	4.23	8.29	5.0	7.52	3.91	2.0	4.36
Flexibility (when some control possible)									
No Input	43%			59%			45%		
At most Input	40%			26%			37%		
Some Control	13%			11%			13%		
A lot of Control	4%			4%			5%		
Advanced Notice (when schedule varies)									
$\leq$ 1 Week	53%			30%			43%		
1–2 Weeks	21%			34%			25%		
3+ Weeks	15%			8%			15%		
Job Satisfaction (1-Like Very Much $\rightarrow$ 5-Dislike very much)									
Job Satisfaction	2.2	2.0	1.07	2.22	2.0	1.15	2.24	2.0	1.14
Job Satisfaction (IR>0.34)	2.53	2.0	0.99				2.6	2.0	1.02
Sample Size		219			186			405	

Selected summary statistics for the primary sample of retail jobs reported in the NLSY97. The third column uses the nationally representative weights. All monetary values are in 2013 dollars. The *instability ratio* provides a measure variation in week-to-week hours and is defined as the range of hours divided by usual hours.

Table 4 contains descriptive statistics for the sample of jobs. Around fifty-four percent of the observed jobs report hours that vary from week to week. This moment is robust to re-weighting the sample using the nationally representative weights provided by the NLSY97 as seen in the third column. Figure 1 plots the variability in hours across the spectrum of usual hours but a simpler measure of variability in depicted in Figure 5 for workers whose hours vary. The horizonal axis denotes the range in hours the respondent reported over the last four weeks while the vertical axis records the frequency of each answer. This figure reiterates that the sample's median range of over eight hours implies that around half of workers face variation equivalent to a full shift or more in hours. Finally, the instability ratio (IR) from Lambert et al. (2014) provides a useful measure of the extent of variability in hours conditional on mean hours. The IR ratio divides the range of hours by usual hours to capture the principle that a given range of fluctuation in hours has a larger impact on jobs where mean hours are smaller. To illustrate, conditional on hours that vary the mean IR is 0.34 but jumps to 0.58 for part-time workers demonstrating that absolute deviations may understate the degree of week-to-week variation.

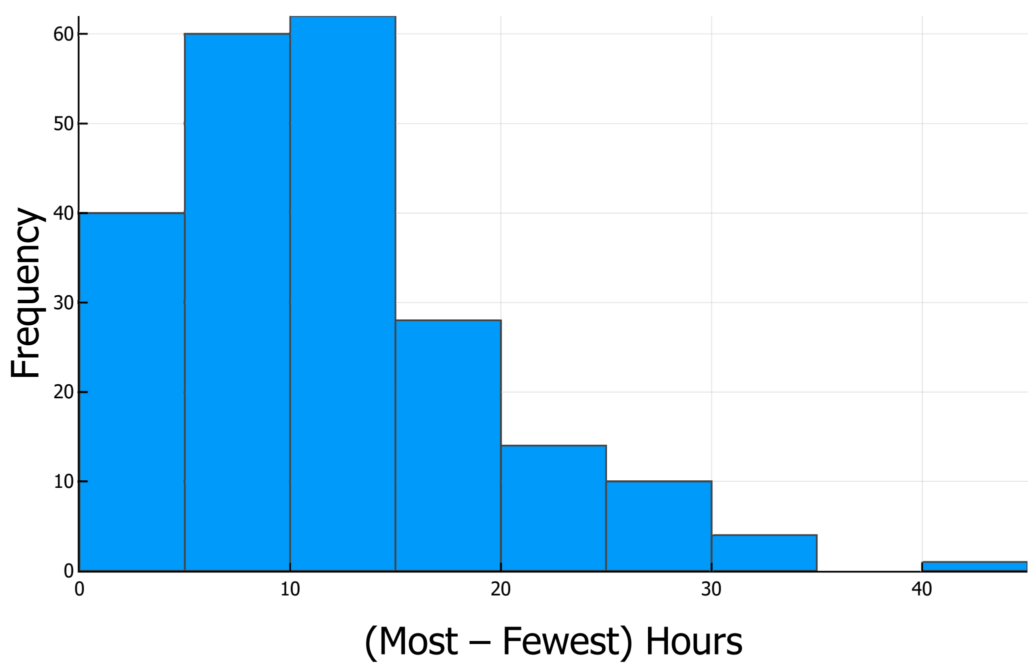

Figure 5: The distribution of the range in weekly hours for the sample of retail employees.

The average job in the sample pays around twice the minimum wage at $14 an hour for around thirty-five hours per week of work. The similarity between mean hours and wages for fixed and variable schedules may hide relationships suggested by some simple correlations. For example, the OLS coefficient from a regression of wage on IR has a statistically significant coefficient of -3.27 hinting at the importance of accounting for higher order moments when describing a job. Around fifty-seven percent of the sample reports usual hours between thirty-five and forty hours underscoring that these jobs are not necessarily entry level or minimum wage type jobs and that varying hours and lack of control over schedule are features of the wider retail labor market. Recall that since sample selection conditions on labor force participation and age profile the average respondent is around thirty years old which puts into perspective that fifty-seven percent of respondents have at least one child (for more, see Table 5).

VI. Estimation of Preferences

The survey instrument described in Section IV.A collects background information on respondents who are then asked to choose between two job offers and an unemployment benefit on a series of fourteen vignettes. Thirteen of the fourteen vignettes they see are randomly selected from a pool of fifty-one. Figure 3 provides an example of one such vignette. All respondents saw the same first vignette with two jobs meant to resemble entry level positions and the lowest value of unemployment benefit. The sample used here to estimate worker preferences took the survey over the internet in January 2016. All respondents were selected members of a panel maintained by Qualtrics, a private survey institute. Potential respondents were screened by gender and for a birth year between 1980 and 1984 so that the resulting sample had a roughly even gender split and matched the age profile of the NLSY97. The data contains the stated preferences over the fourteen vignettes as well as information on demographics and the job characteristics of employed respondents. These auxiliary questions were intentionally similar if not identical to those asked by the NLSY97 to maximize comparability. Respondents in the final sample were verified using several quality controls in addition to the monitoring and panel maintenance performed by Qualtrics.

To detect inattentiveness on the part of respondents and improve the quality of responses the survey mechanism employs three methods. To be included in the sample respondents must have an IP address located in the United States, take no less than twenty-five percent of the mean total time to finish the survey, and pass an attention filter described in Figure 15 in Appendix B. These measures screened out around sixty percent of potential respondents with the vast majority failing the attention filter. In fact, around sixty-five percent of male and fifty-two percent of female respondents who began the survey failed the attention filter. The remaining 1,026respondents are described in Table 5 which presents selected statistics of the sample as compared to the unweighted and weighted NLSY97 sample of retail workers from the previous section.

The NLSY97 over-samples minorities to improve the quality of their data and provides nationally representative weights that allow extrapolation from their sample to the United States. However, to be included here, NLSY97 respondents must report a primary employer in the retail sector which introduces a

Table 5: UWSF Sample Statistics

	Vignette (unweighted)		NLSY97 (unweighted)		NLSY97 (weighted)	
	Females	Males	Females	Males	Females	Males
General						
% of Sample	50%	50%	54%	46%	52%	48%
Mean Age	33.1	32.9	31.0	30.8	31.0	30.8
Cohabitating	67%	61%	49%	44%	56%	43%
Student	8%	10%	9%	6%	7%	7%
Avg # Jobs	0.9	1.2	1.1	1.1	1.1	1.1
Any Kids	67%	50%	68%	43%	63%	42%
Avg # Kids	1.5	1.0	1.5	0.9	1.3	0.8
Race and Ethnicity						
White	85%	85%	56%	61%	74%	75%
Black	9%	9%	25%	24%	13%	15%
Hispanic	10%	10%	25%	25%	13%	16%
Education: Highest Degree						
< High School	3%	3%	11%	7%	9%	7%
High School	16%	16%	66%	67%	62%	65%
A.A. Degree	12%	12%	6%	6%	7%	7%
B.A.	27%	27%	15%	18%	18%	20%
> B.A.	13%	13%	3%	1%	4%	1%
N	518	508	218	187	218	187

A comparison of selected demographic statistics for the USWF and NLSY97 samples as well as the NLSY97 with nationally representative weights.

non-trivial sample selection that is not corrected in the weighted statistics and as such should be interpreted accordingly. Comparatively, the UWSF respondents are slightly older, more highly educated, and more likely to report their race as white or Caucasian. Additional summary statistics are available in Tables 11 and 12 in Appendix B which show that around seventy-two percent of respondents report a primary job of which a third are part-time with instability ratios higher than found in primary sample of jobs. However, a degree of caution is warranted as, for example, respondents were mixed on whether their time spent taking surveys qualified as employment.

Other than matching the age profile and residing in the U.S., the sampling procedure used for the UWSF survey does little to ensure a similar sample to that in the NLSY97. Even if the panel of potential respondents was nationally representative, the over-sampling of minorities and selection on retail jobs makes it essentially infeasible to replicate in a cost efficient manner since the NLSY97 respondents come from across the United States and from multiple retailers.[11] To help mitigate differences between the samples I weight observations from the UWSF sample in the first stage estimation of preferences to match observed proportions in the NLSY97 sample. To create the weights, I choose four demographic characteristics available in both samples that hopefully capture major sources of worker heterogeneity and correct for differences in sample composition so that the estimated preferences more closely resemble what you find if you gave the survey to the NLSY97 sample. The characteristics used are binary indicators for gender, non-white, having a B.A. degree or higher, and the presence of a child under age six. This creates sixteen bins where the frequency in the NLSY97 of a given bin determines the weight on the contribution of the relevant bin in the vignette sample. See Table 6 for a summary of the characteristics in both samples.

I interpret the first three characteristics as plausibly exogenous conditions from the perspective of individuals in their late-twenties or early-thirties choosing between jobs. The fourth, having a young child, stands apart since there is likely an (un-modelled) endogenous relationship between labor participation and fertility decisions. I consider incorporating the fertility decision beyond the scope of this paper as it adds little to the proposed counterfactuals after accounting for consequent heterogeneity and requires a separate identification strategy for preferences over children. However, in recognition of the presumed importance the presence of a young child has over preferences for the kind scheduling instability found in the jobs of this model, I attempt to allow the unexplained heterogeneity to capture this effect by weighting the sample accordingly. No bins are empty from either sample. Estimated preferences reflect the re-weighted UWSF sample.

To comment on the reliability of the discrete experiment approach, consider the well-known critique of

[11]One possibility would be to draw the entire dataset from a large single retailer where employees take both the vignette survey and report job characteristics. This approach helps mitigate the concern of heterogeneity in the structural parameters underlying firm technology but at a cost of making counterfactuals specific to that retailer and its employees.

Table 6: Statistics from Variables Used to Re-weight Preferences

	UWSF		NLSY97		NLSY97 Wtd.	
	Females	Males	Females	Males	Females	Males
Children Under the Age of 6						
White & ≥ B.A.	11%	12%	3%	5%	4%	6%
Non-White & ≥ B.A.	2%	3%	2%	2%	2%	1%
White & < B.A.	12%	11%	19%	12%	23%	14%
Non-White & < B.A.	3%	2%	14%	12%	7%	8%
No Children Under the Age of 6						
White & ≥ B.A.	21%	24%	9%	10%	12%	12%
Non-White & ≥ B.A.	4%	4%	4%	3%	3%	2%
White & < B.A.	41%	36%	25%	34%	34%	43%
Non-White & < B.A.	6%	8%	24%	22%	14%	14%
% of Sample	50%	50%	54%	46%	52%	48%

The UWSF sample was re-weighted for estimation along four demographics characteristics to more closely match the sample of retail jobs. $\geq$ *B.A.* indicates the respondent has a B.A. degree or higher. This table compares the moments of the weighting variable to the un-weighted sample and nationally weighted NLSY97 sample. The unweighted NLSY97 sample is used for second stage estimation.

stated preference data in Diamond and Hausman (1994) who observe that when considering hypothetical options there is no guarantee that respondents have preferences over all options or respect budget constraints when providing answers. Their analysis primarily considers surveys soliciting a willingness to pay value in dollars for a public good and outline serious concerns for extrapolating from this type of data. More recently, the literature has adopted the vignette-style discrete choice experiment used here that measures willingness to pay using a revealed preferences type argument and potentially reduces the noise inherent in soliciting point estimates. Nevertheless it remains to be shown that the constructed jobs do not have attributes that would be unfamiliar or irrelevant to the respondents. Table 3 describes the set of possible attributes for vignette jobs. An examination of the reported characteristics of NLSY97 jobs in Table 4 confirms that the attributes are well-represented in equilibrium jobs and, more directly, Table 11 summarizes the characteristics of jobs reported by the UWSF sample which similarly suggests that respondents would find the vignette attributes realistic independently if not necessarily in every particular combination.[12] Together, this evidence suggests that the non-wage attributes of vignette jobs were plausible draws from the set of equilibrium jobs and that the UWSF sample would have experience and preferences over the set of attributes.

The object of estimation is to recover the structural parameters underlying preferences, the extent of heterogeneity in population, and the stigma cost of unemployment. As respondents compare the hypothetical offers, correlation in observed choices across individuals reveals common preferences. Variation in choices

[12]Very few respondents report 3–4 weeks of advanced notice in either the NLSY97 or the UWSF sample. To maximize statistical power the survey instrument reduces the set of options to $\leq$ 1 week, 1–2 weeks, and 3 or more weeks.

across individuals recovers the heterogeneity in preferences which, in line with the model, will be captured in the risk aversion parameters $(\rho_c, \rho_\ell) \sim F_\rho$. The choice between the job offers and the unemployment benefit identifies the participation decision of workers and, in particular, any reluctance to choose employment despite a relatively attractive value of unemployment recovers the stigma cost of unemployment. In the model the stigma cost depends on a set of demographic characteristics that employers are legally prohibited from using to discriminate in labor contracting. However, since the exact relationship between the stigma cost and these characteristics is neither of first order importance for the dynamics of the model nor something that the survey was designed to reveal, I make the simplifying assumption in estimation that the stigma cost is $g(z) = \overline{g}$, or a fixed dis-utility independent of worker characteristics.

To map the worker's preferences over consumption and hours worked to the jobs in the vignette let the utility of a vignette job be

$$Z(C, L, X; \rho) = \frac{C^{1-\rho_c} - 1}{1 - \rho_c} - \delta \frac{\ell^{1-\rho_\ell} - 1}{1 - \rho_\ell} + \psi X$$

where X contains indicators for the level of advanced notice and flexibility. The parameter ψ captures the valuation of these secondary job attributes which represent incidental parameters to the larger paper. Analogously, the utility of the unemployment option on a vignette is

$$Z(C, 0, 0; \rho) = \frac{C^{1-\rho_c} - 1}{1 - \rho_c} - \overline{g}.$$

When choosing between alternatives on a vignette, respondents make decisions without knowing the actual sequence of future pay-offs so the value of alternative k for person i on vignette j is

$$V_{ijk} = \mathbb{E}\left[Z(C_{ijk}, L_{jk}, X_{jk}; \rho)\right] + \nu_{ijk}$$

where ν_{ijk} is an unobserved choice shock specific to individual, vignette, and alternative. Since the wages for each choice are adapted to the individual, C_{ijk}, varies by individual but hours, L_{jk} and secondary attributes, X_{jk}, are known given the vignette and alternative. I assume that the distribution of F_ρ is Log Normal with an unknown mean and covariance matrix so that

$$\begin{pmatrix} \ln \rho_c \\ \ln \rho_\ell \end{pmatrix} \sim N\left(\begin{bmatrix} \rho_{c0} \\ \rho_{\ell 0} \end{bmatrix}, \begin{bmatrix} \rho_{c1}^2 & \cdot \\ \rho_{c\ell}\rho_{c1}\rho_{\ell 1} & \rho_{\ell 1}^2 \end{bmatrix} \right).$$

The logarithm forces ρ to reside in the set $(0, \infty)$ as is consistent with the constant relative risk aversion form of utility.

For a respondent i to choose alternative k from the set of three possibilities on each vignette j it must satisfy

$$V_{ijk} = \max\{V_{ij1}, V_{ij2}, V_{ij3}\} \tag{28}$$

where the first two options are job offers and the third option is an unemployment benefit. As econometricians, we observe the full set of characteristics of all alternatives and the functional forms of $Z(\cdot)$ up to a set of parameters but not the individual's $\rho = \{\rho_c, \rho_\ell\}$. Given ρ, if we assume that ν follows a Type 1 Extreme Value distribution then the difference $\nu_{ijk} - \nu_{ija}$ has a logistic distribution and the solution to (28) is probabilistic. These assumptions motivate the mixed Logit model where the unobserved but fixed and individual-specific ρ allows for arbitrary correlation across alternatives for individual i while mapping to the heterogeneity of workers in the model. Let L_{ij} be the probability of person i's observed choice on vignette j and collect all observables into W. Define θ as the preference parameters excepting ρ. An individual's contribution to the likelihood is their sequence of observed choices across J vignettes, or

$$\mathbb{L}_i(\theta, \rho) = \prod_{j=1}^{J} L_{ij}(\theta, \rho; W).$$

To estimate the model I use simulated maximum likelihood estimator and integrate out the unobserved ρ. Formally, let Θ be the set of preference parameters θ and the mean and variance parameters $(m\rho, s\rho)$, then the estimator finds Θ to maximize the simulated log-likelihood

$$L(\Theta; X) = \sum_{i=1}^{N} \log \left(\sum_{s=1}^{S} \prod_{j=1}^{J} L_{ij}(\theta, \rho^s(m\rho, s\rho); W) \right), \tag{29}$$

where S is the number of simulations used in Monte Carlo integration.

Table 7 presents the results of estimation. All parameters are statistically significant at the five percent level. Table 14 reports the similar results obtained when estimating without re-weighting observations. The weight on labor hours relative to consumption is 0.65 which is relatively small given consumption, $w\ell$, has a larger scale than labor hours, ℓ. The parameters governing the joint distribution of the coefficients of relative risk aversion are perhaps most easily interpreted through the bivariate density plotted in Figure 6. The slight right skewness in both sources of risk aversion as well as the apparent correlation between them highlight the importance of considering the full distribution over point estimates to explain observed labor market behavior. Overall, the absence of heterogeneity is rejected by a Wald test which underscores the importance of accounting for this heterogeneity when considering the welfare costs of policy for the labor force at large.

The coefficient of risk aversion for consumption ranges from 0.66 to 1.04 between the tenth and ninetieth

Table 7: Estimated Preference Parameters

Parameterization

$$\frac{C^{1-\rho_{c,i}}-1}{1-\rho_{c,i}} - \delta\frac{L^{1-\rho_{l,i}}-1}{1-\rho_{l,i}} + \psi X - \mathbb{1}_{\text{ue}}\overline{g} \qquad \begin{pmatrix} \ln\rho_c \\ \ln\rho_\ell \end{pmatrix} \sim N\left(\begin{pmatrix} \rho_{c0} \\ \rho_{\ell 0} \end{pmatrix}, \begin{bmatrix} \rho_{c1}^2 & \cdot \\ \rho_{c\ell}\rho_{c1}\rho_{\ell 1} & \rho_{\ell 1}^2 \end{bmatrix} \right)$$

Re-weighted Estimates

	δ	ρ_{c0}	$\rho_{\ell 0}$	ρ_{c1}	$\rho_{\ell 1}$	$\rho_{c\ell}$	$\overline{g}$
Estimate	0.65	-0.20	0.56	0.21	0.50	0.58	0.79
SE	(0.19)	(0.01)	(0.14)	(0.01)	(0.09)	(0.01)	(0.16)

	Advanced Notice		*Flexibility*		
ψ :	1–2 Weeks	3+ weeks	A little	Some	A lot
Estimate	0.26	0.31	0.54	0.82	1.23
SE	(0.04)	(0.04)	(0.05)	(0.05)	(0.04)

N=1026 Log-Likelihood: -927.7

Estimates for the structural parameters underlying worker preferences and heterogeneity. Parameters were estimated using simulated maximum likelihood on data from a purpose-made discrete choice experiment within the UWSF survey. Observations were weighted by four demographic characteristics to more closely matched the NLSY97 sample of jobs. Simulation used 1,000 Halton draws and standard errors were calculated using the outer-product method.

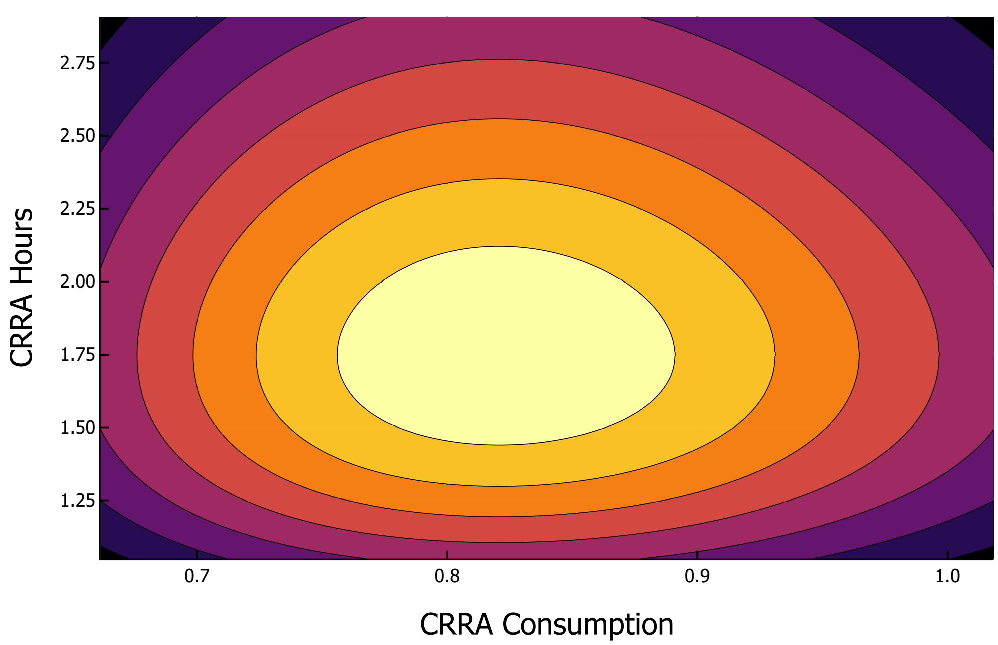

Figure 6: A heat map illustrating the estimated distribution of the risk aversion parameters for consumption and leisure in the population. Values reflect the re-weighted UWSF sample.

percentiles with a median value 0.82. Figure 7 plots the marginal distribution for the coefficients of relative risk aversion. These estimates come from data that is demographically more diverse and sit in between the estimates typically found in the micro literature (lower) and those from the macro literature. To maximize comparability I look at a few studies that also have a model involving decisions with long term uncertainty. For example, Joubert and Todd (2016) find a coefficient of 0.42 when estimating a labor participation and pension funding model and Navarro and Zhou (2017) report a value of 0.49 when estimating a life-cycle model of labor participation and schooling choices under credit constraints. The range I find is slightly lower than the typical calibration values in the macro literature (see Thimme (2017)). One important caveat when comparing these values with the literature is that the model and estimation method both impose that observed hours are not necessarily optimally chosen hours such that we would observe different consumption and leisure choices if, like in the neoclassical labor supply model, workers were allowed to adjust their labor supply given the wage. In some respects, the set-up here is perhaps more comparable to Metrick (1995) who uses the betting decisions of players on Jeopardy! to estimate a risk aversion parameter of 1.02.[13]

The coefficient of risk aversion for labor hours ranges from 1.12 to 2.83 with a median value of 1.72 indicating a right tail skewness in the marginal distribution. Figure 8 plots the marginal distribution for the risk aversion over labor hours. These values indicate a uniformly higher degree risk aversion over hours relative to consumption. The findings may partially reflect the extent of variation in hours of plus or minus fifty percent of mean hours used in the vignettes to capture preferences over a large range in variability. Here there is also an issue that both consumption and hours vary when hours vary but the comparatively smaller variation in the wage rate of plus or minus ten percent left less independent variation in consumption. For comparison, Navarro and Zhou (2017) estimate the coefficient of relative risk aversion for labor of around 1.0. These statistics are also commonly reported as Frisch elasticities which measure the response in hours to a change in the wage rate and here is equivalent to $1/\rho_\ell$. The estimated median Frisch elasticity of 0.58 sits within the set of typical estimates from the micro literature though the upper bound is closer to the typical values of one to two common in the macro literature as discussed in Keane and Rogerson (2012).

To provide some evidence of the reliability of the preferences estimated here I find it useful to consider the implied willingness to pay. Consider an individual working thirty hours a week on average but facing week-to-week variation amounting to plus or minus eight hours across weeks. Figure 9 plots the percentage wage reduction that individual would accept to move to a job with exactly thirty hours each week. The median willingness to pay is around two percent. This point estimate coincides with the fact that around eighty percent of individuals would not accept more than a six percent reduction but obscures the fact that

[13]Cohen and Einav (2007) translate the absolute risk aversion parameter from Metrick (1995) to its relative risk aversion equivalent.

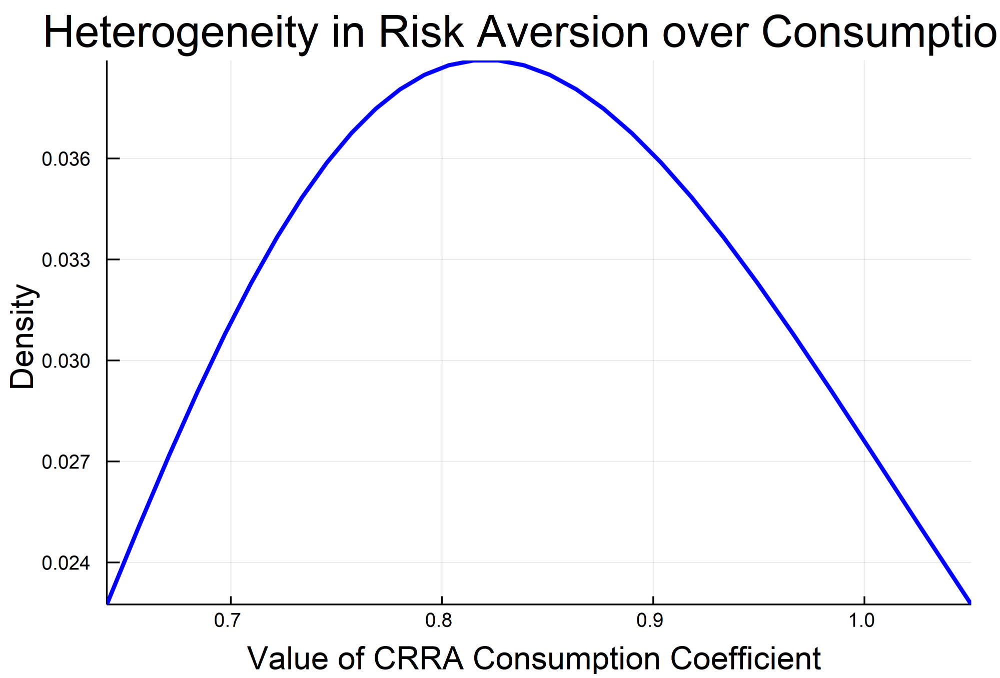

Figure 7: The marginal distribution for the estimated heterogeneity in the coefficient of relative risk aversion for consumption. The values were estimated using a sample from the UWSF survey.

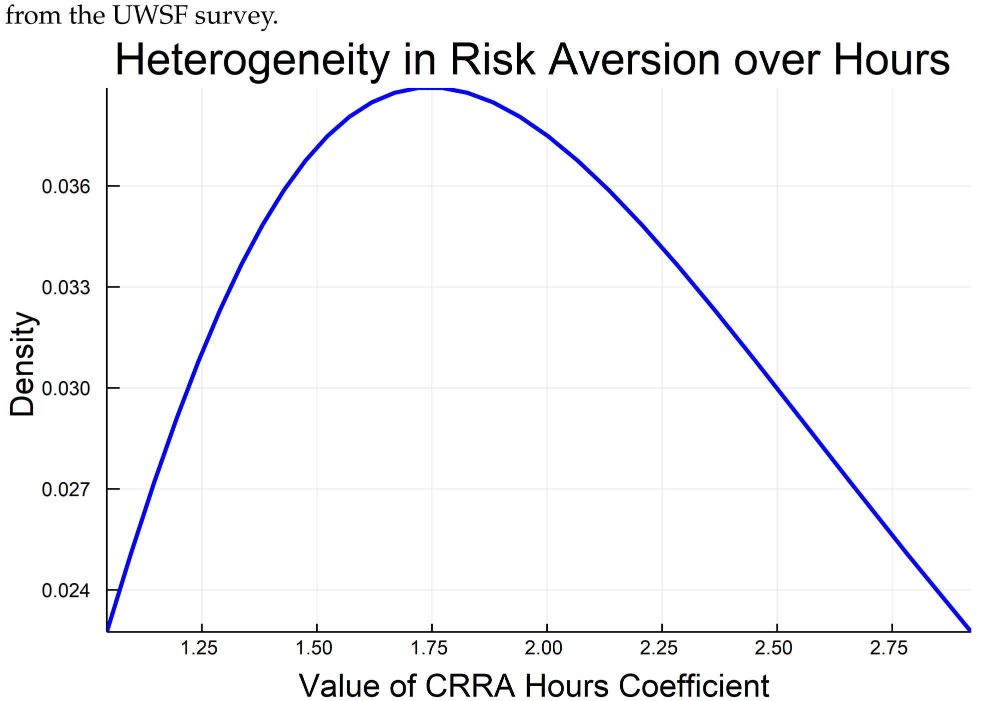

Figure 8: The marginal distribution for the estimated heterogeneity in the coefficient of relative risk aversion for labor hours. The values were estimated using a sample from the UWSF survey.

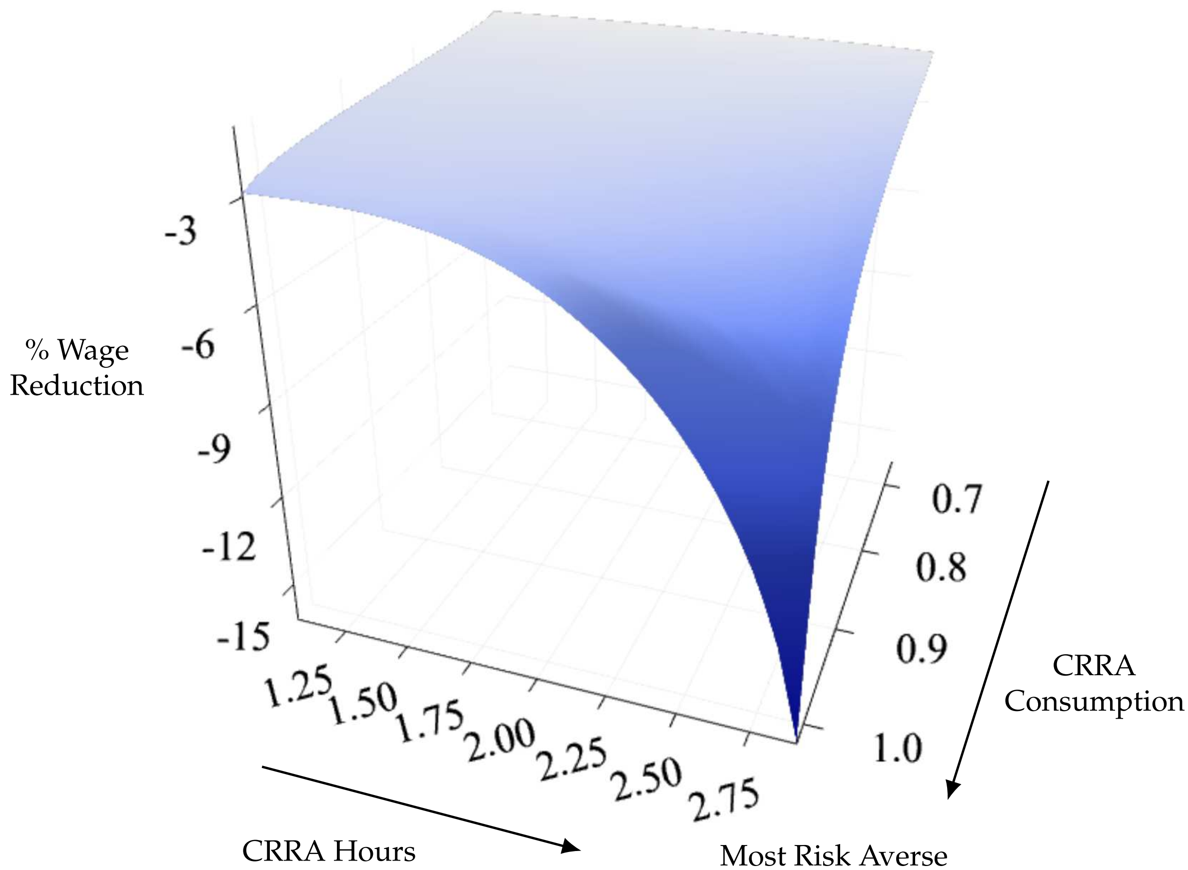

Figure 9: The estimated willingness to pay of individuals as the percentage wage reduction they would accept to move from a thirty hours per week job with plus or minus eight hours variation week to week to one with exactly thirty hours each week. Based on the preferences estimated using the UWSF survey.

a small portion of workers would give up as much as sixteen percent to move to the fixed schedule. Here we intuitively see a group that might benefit from a policy that limits their variation in hours even if most workers are relatively indifferent.

The stigma cost of unemployment is perhaps most easily interpreted when translated into dollar values. An individual with the median risk aversion parameter who receives a three hundred dollar per week unemployment benefit has the utility equivalent to a ninety-five dollar loss of income from the stigma. The skewness of the risk aversion parameter results in a comparable skew of stigma where the tenth and ninetieth percentiles are around thirty-five and five hundred dollars, respectively.

The estimated preference parameters of interest are robust to changing which secondary attributes are included in X. Estimating the model with ψ_i where each secondary attribute has an unobserved person-specific value drawn from a joint Normal distribution that is entirely independent of ρ improves model fit but has little effect on the primary parameter estimates. Since the coefficients on advanced notice and flexibility are not of immediate interest to this paper I will limit discussion of the results. The estimates are increasing in the intuitive direction and seem plausible given the nature of the vignette jobs and related literature (see Eriksson and Kristensen (2014) and Mas and Pallais (2016)). For a job entailing 35 hours a week at $12.50 an hour, the median individual would accept a two and a half percentage point wage reduction to move from one week or less of advanced notice to at least two weeks. For a similar set up, a person will forgo six percent of their wage to have some input over the timing of their hours and a twenty-seven percent reduction to have some direct control over when they start and stop working. In both cases, once a job has some advanced notice or some flexibility the returns to additional degrees of control are substantially smaller than the initial move from no control to some control.

VII. Estimation of General Equilibrium Model

To estimate the remaining structural parameters relating to firm and search technology I solve and simulate the equilibrium model of the labor market taking as given worker preferences and heterogeneity estimated in the first stage. The solution generates a set of equilibrium labor contracts from which I draw cross-sectional moments on wages, hours, and employment flow statistics that I fit to comparable moments from the retail sector. Table 8 summarizes the seventy-six moments used in estimation. While the observations from the first stage estimation were weighted to match the demographic characteristics of the NLSY97 sample of retail jobs, nothing fundamental to the UWSF survey ties the respondents to the retail sector other than the characteristics of the jobs offered in the discrete choice experiment. In the second stage of estimation the moments are specific to the retail sector whenever possible and so the technology, heterogeneity in productivity shocks, and search dynamics should be interpreted accordingly.

The primary structural parameters underlying firm technology are the scale and labor elasticity of the production function. These parameters are identified by matching the moments of hours and wages in observed and simulated jobs. In practice, given a set of parameters the model generates optimal labor contracts for all observed firm-worker pairings which in turn implies a wage and a distribution of weekly hours. The common unemployment benefit factors into setting the level of compensation and is therefore helpful in matching observed wages. I assume the distribution of firm types is Log Normal and known up the mean and scale parameters such that $\sigma \sim \ln N(m_\sigma, s_\sigma)$. The distribution of firm types is recovered from the observed variation in hours. The observed distribution of hours within a contract stems from the choices firms make given the wage and the sequence of shocks they faced. Therefore, the variation in hours relates directly to the firm's type, the variance of their productivity shocks, as can seen in the following identity

$$\mathrm{Var}(\ln \ell) = \frac{\sigma_n^2}{(1-\alpha)^2}.$$

Variation in hours across contracts contains information about the distribution of firm types such that even though the variance of the shocks specific to any particular pairing is not identified, the distribution of such firm types is recoverable.

The equilibrium search dynamics are governed by the set of parameters specifying the costs of creating and maintaining a vacancy, the matching technology, and the exogenous rate of job destruction. Though treated here as structural parameters, the matching technology is better understood as a reduced form approximation of how search produces meetings between potential pairs. The cost of creating a vacancy directs the sorting of workers between the fixed and variable sectors by controlling firm entry. Mechanically, as firms enter the probability of a worker contacting a vacancy increases which attracts workers. Since firms are free to enter if they anticipate a positive return, entry occurs until the free entry condition is satisfied. The scope for a positive return varies directly with the cost of creating a vacancy. An increase in the costs of creating a vacancy, c_E and c_v, reduces the firm's expected value of entering the sector for any given set of participating workers and thus helps determine whether a particular $\{\sigma, \rho\}$ pairing is sufficiently valuable for the firm to accept the match. Non-viable pairings determine the values of $\eta_m(\rho)$ and the division of the workers between the sectors.

I assume the costs of maintaining a vacancy $\chi_f(\theta_f)$ and $\chi_v(\theta_v)$ are linear in the labor market tightness of the sector where $\chi_m(\theta_m) = \chi_m^1 \theta_m$. To allow for costs of filling a vacancy to increases as the demand for labor grows I assume $\chi_m^1 > 0$ such that the costs of maintaining a vacancy are increasing in labor market tightness. The movement of workers from one sector to another has consequences for the profitability of participating firms and, in practice, the maintenance costs help fine tune the equilibrium value of θ_m and thus the measure of vacancies created in sector m. Though the cost parameters do influence labor contracts through the wage setting equation, the salient moments for the estimation of the cost parameters are the ratio of contract types, aggregate labor market tightness, the aggregate unemployment rate, and the rate of job openings which all depend directly on the job and worker contacts rates. These rates, in turn, are functions of the sector specific labor market tightness and related costs.

The fixed and variable sectors operate independently but use the same technology which for a Cobb-Douglas matching function with constant returns to scale is fully summarized by the scale and the elasticity of matches with respect to unemployment. Since not all meetings result in viable matches the estimation

Table 8: List of Moments Used for Second Stage Estimation

Description	Statistic	Comments
Search Dynamics		
Unemployment	Rate	Aggregated portion of unmatched workers at the end of a period. Set at the aggregate unemployment rate in the U.S. economy
Labor Market Tightness	Ratio	Aggregated ratio of job openings (vacancies) to unemployed (unmatched) workers. Set to the total number of job openings (all industries) in JOLTS divided by the level of unemployed workers
Job Openings	Rate	Aggregated measure of job openings (vacancies). Set to the total number of job openings in the retail sector from JOLTS. Equal to the number of job openings divide by the sum of employment and job openings.
$\frac{\text{variable jobs}}{\text{fixed jobs}}$	Ratio	Ratio of jobs in variable market to fixed market. Set using ratio in observed jobs in NLSY97 sample.
Wage and Hour Moments from Fixed Sector		
Wages	Quantile	Taken from NLSY97 sample of jobs.
Hours	Quantile	Around 60% of observations are 40 hours per week. Taken from NLSY97 sample of jobs.
$\frac{\text{Wages}}{\text{Hours}}$	Quantile	Captures correlation between wages and hours across jobs. Taken from NLSY97 sample of jobs.
Wage and Hour Moments from Variable Sector		
Wages	Quantile	Taken from NLSY97 sample of jobs.
Hours	Quantile	Mean of distribution of actual hours. Taken from NLSY97 sample of jobs.
Variance of Hours	Quantile	Variance of distribution of actual hours. Taken from NLSY97 sample of jobs.
$\frac{\text{Wages}}{\text{Hours}}$	Quantile	Captures correlation between wages and hours across jobs. Taken from NLSY97 sample of jobs.
$\frac{\text{Variance of Hours}}{\text{Hours}}$	Quantile	Captures correlation between mean and variance of the distribution of actual hours across jobs. Taken from NLSY97 sample of jobs.

A description of the 76 moments used in estimation. All quantiles are the tenth to ninetieth deciles.

procedure entails solving for model equilibrium which directly enforces selection into the simulated moments. These parameters will largely be identified by how they affect the vacancy and worker contact rates for a given labor market tightness which in turn determine the unemployment and job opening rates.

One useful consequence of the steady state equilibrium is that since employment inflows must match outflows I can confidently calibrate the exogenous rate of job destruction, s, to the rate observed in the data. To measure inflow I take the average of the JOLTS non-seasonally adjusted monthly rate of hires in the retail industry between November 2013 and July 2014. The resultant value of 4.7 has a sample variance of 0.08 over the same period. Ultimately the choice to use the reported hires over the equally valid series measuring total separations (with a mean rate of 4.6) is motivated by the endogenous hiring decision in the model and the absence of endogenous firing decisions. I plot both series in Figure 16 which shows that while not exactly coincident the series do appear closely related.

The estimation concept is a simulated method of moments estimator where for some set, g^d, of moments calculated from the equilibrium jobs and worker flows there is comparable mapping, g^s, from the structural parameters listed above to analogous moments from a simulated equilibrium. Let Ω be the set of simulated method of moments parameters and Γ be the preferences estimated in the first stage. $g^s(\Omega;\widehat{\Gamma})$ is the mapping of the structural parameters into simulated moments taking as given the first stage estimates. As a result, the estimated structural parameters minimize

$$\min_{\Omega}(g^d - g^s(\Omega;\widehat{\Gamma}))^T\Sigma(g^d - g^s(\Omega;\widehat{\Gamma}))$$

where Σ is the weighting matrix. When calculating the limiting distributions of the second stage estimates $\widehat{\Omega}$ I ignore any additional error introduced by using the first stage estimates $\widehat{\Gamma}$.

My approach to estimation involves solving for the model equilibrium for a given set of parameters and then simulating the data that generates the moments above. Specifically, I discretize the set of worker and firm types and solve for the steady state equilibrium fully characterized by (θ_v, θ_f) using the algorithm in Appendix C. These sector-specific θ_m and resultant acceptance rates $\eta_m(\rho)$ imply a division of workers between the sectors as a result of their optimizing behavior and directed search. The equation for the steady state value of $u_m(\rho)$ defined in (20) also fully determines the distribution of unmatched workers in sector m, F^m_ρ, the set of vacancies $v_m = \theta_m u_m$, and thus the distribution of steady state matches H_m. To simulate the moments I draw a large sample of jobs from H_m for both sectors from which I can calculate the quantiles of wage and hours and the ratio of variable to fixed jobs. The remaining moments have specific formulas that aggregate up from sector-specific rates like u_m, v_m, and the implied contact rates. Given the moments I can evaluate the loss function and search over the convex parameter space until convergence.

The efficient simulated method of moments estimator uses the inverse variance-covariance matrix of the moments as a weighting matrix. In this application, computing the variance-covariance matrix has proven problematic and sensitive to the underlying approach and assumptions. The moments related to the 405 equilibrium jobs are quantiles of the distributions related to hours and wages where, since the quantile statistics may not have well-defined limiting distributions, the variance-covariance matrix can be constructed through bootstrapping the sample. However, in practice, the sector-specific moments require the sub-division of the sample which reduces the statistical power when bootstrapping which in turn has implications for estimates and model fit. The workforce flow moments are computed as the mean across the monthly series for the nine months between November 2013 and December 2014 which corresponds to the period over which the NLSY97 interviews occurred. While the sample variances and covariances can be computed for these series, the point estimates rely on nine observations. Moreover, the model implies a correlation between search dynamics and observed wages and hours though there is no straightforward way to calculate such moments.

To minimize the sensitivity of the estimates to the assumptions needed for optimal simulated method of moments I choose a weighting matrix where all moments are of the same order of magnitude. I interpret this approach as trying to match all moments as closely as possible in light of the concern that using the optimal weighting matrix may de-emphasize moments that pin down primitives crucial to analyzing counterfactual policy. Given the assumption of a steady state equilibrium the sampling error for all of the moments should be zero which motivates an attempt to try and match each point estimate as closely as possible. To make the argument more concrete, consider the ratio of variable to fixed jobs where through bootstrapping with replacement the statistic has a standard error such that a ratio of less one is within two standard deviations. The possibility that the optimal estimator generates a simulated economy with fewer fixed jobs than variable raises the issue of whether this equilibrium is an appropriate economy in which to run the kind of counterfactual policy experiment under consideration here.

One feature of the data poses a particular problem for model fit. The set of observed jobs with fixed schedules has a large mass point at forty usual hours per week. This feature is illustrated in Figure 10 which plots the quantiles for mean hours in fixed contracts where around sixty percent of all observations report forty hours as usual. Given the curvature of the production function the only feature in the model that would generate a comparable mass point in the simulated set of jobs is a mass for some particular ρ value that participates in the fixed sector though this occurrence would further imply a (non-existent) corresponding mass point in wages as well. Either way, the assumptions used to estimate worker heterogeneity in the first stage effectively rule out this possibility. More generally, even if estimation allowed for this possibility, this mechanism does not provide a satisfactory explanation of how this feature arises from the economy's

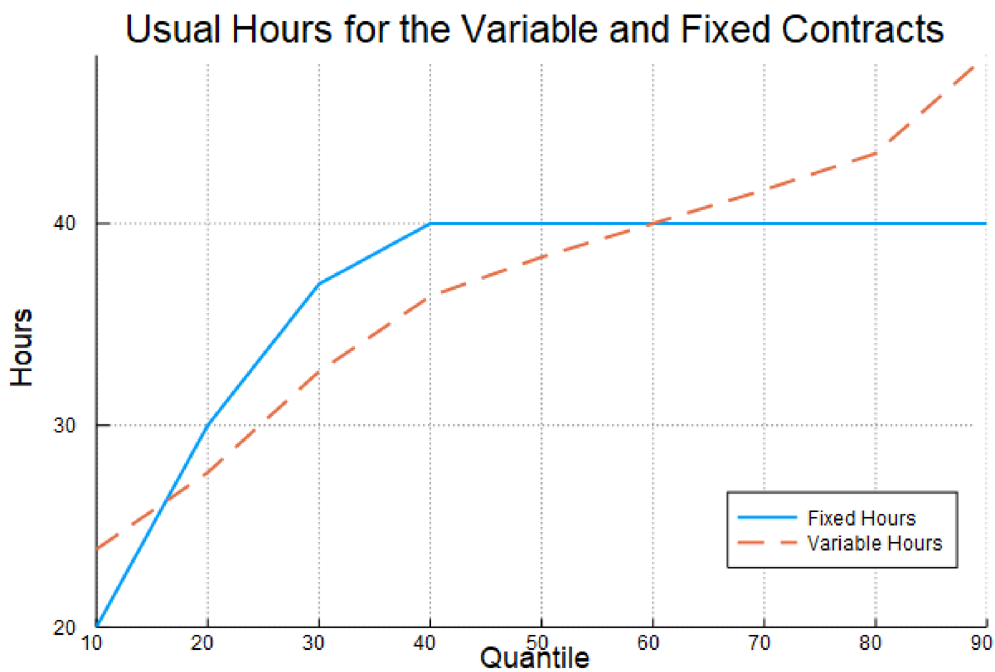

Figure 10: The observed profile of usual hours in equilibrium jobs for variable and fixed contracts. The set of fixed jobs have a mass point of around sixty percent of observations at forty hours per week that the model cannot endogenously replicate.

primitives.

The results from estimation are presented in Table 9. The elasticity of labor demand is relatively elastic at around 0.66 which coincides with the estimate in Cooper et al. (2015) who model a similar production technology that abstracts away from the fixed factors of production and incorporates stochastic productivity but estimated with data from manufacturing plants. The elasticity of the matching function with respect to the unemployment rate is 0.70 which broadly align with what Cooper et al. find using data for the entire U.S. labor force. The scale of the matching function is largely determined by trying to match the unemployment and job opening rates. The value of 0.49 is most easily interpreted by noting that unemployed workers in the variable and fixed sectors both have around an eighty percent chance of contacting a vacancy within four weeks.

The common unemployment benefit is estimated to be around four hundred and fifty dollars per week which is considerably higher than the average U.S. figure of three hundred dollars. I interpret the large value as resulting more from choices in model parameterization than a solid identification of the outside option for retail workers. Within the model the unemployment benefit largely serves to set levels of labor compensation though the stigma cost of unemployment and heterogeneity in worker preferences imply that workers are willing to accept a job with considerably lower amount of per-period compensation. For intuition, setting $\overline{g}$ arbitrarily small results in a much lower estimate of $\widehat{b}$ but generates a much worse model fit. The operative

Table 9: Estimated Firm and Matching Parameters

Parameterization

$$y = e^{\varepsilon}\lambda\ell^{\alpha}$$
$$\sigma_n \sim \log N(m\sigma, s\sigma)$$
$$M = \mu u^{\tau} v^{1-\tau}$$

Firm Technology and Heterogeneity

	α	λ	$m\sigma$	$s\sigma$	b	
Estimate	0.66	94.63	-3.41	0.05	467.53	
SE	(0.05)	(36.84)	(0.63)	(0.01)	(13.24)	

Search

	c_E	c_v	χ_f^1	χ_v^1	τ	μ
Estimate	3976.74	12187.36	731.08	187.78	0.70	0.67
SE	(1841.17)	(3903.86)	(16.45)	(26.44)	(0.03)	(0.06)

Number of Moments = 76 — Minimum Value: 98.84

Estimates for the structural parameters underlying firm technology, firm heterogeneity, and aggregate matching. Parameters were estimated using Simulated Method of Moments where the structural parameters were used to solve for the model equilibrium and simulate a series of wage and hours profiles as well as aggregate moments with observable analogues in data from the retail sector.

mechanism is that ρ amplifies the dis-utility $\overline{g}$ which appears in the net value of employment term within the participation constraint (equation (8)) which binds in contract bargaining. The direct implication is that more risk averse workers will accept a lower amount of labor compensation all else equal. Therefore, I view the estimated unemployment benefit as closer to a model fitting parameter than a minimum value of sustenance and identify it as an area for improvement—possibly through giving workers a positive bargaining weight. Figure 11 plots the implied distribution of firm types where the median value is close to zero but the long right skew implies that a small percentage of firms see considerable variability in productivity shocks. For interpretation, values of $\sigma < 0.05$ correspond to a range in hours smaller than five across weeks.

To interpret the costs of creating and maintaining a vacancy it is helpful to remember that firms are infinitely lived and face a low probability of match destruction. Therefore, the entry cost of around four thousand dollars is the present value of receiving around eighty dollars each week in profit forever. The much higher added cost of entering the variable sector is equivalent to an extra two hundred and forty dollars per week and partially reflects the premium a high variance firm can make by adjusting hours each period. The ranking is reversed for the costs of maintaining a vacancy where, once adjusted given the labor market tightness ($\chi_m(\theta_m) = \chi_m^1\theta_m$), the cost in the fixed sector amounts to two hundred dollars per week

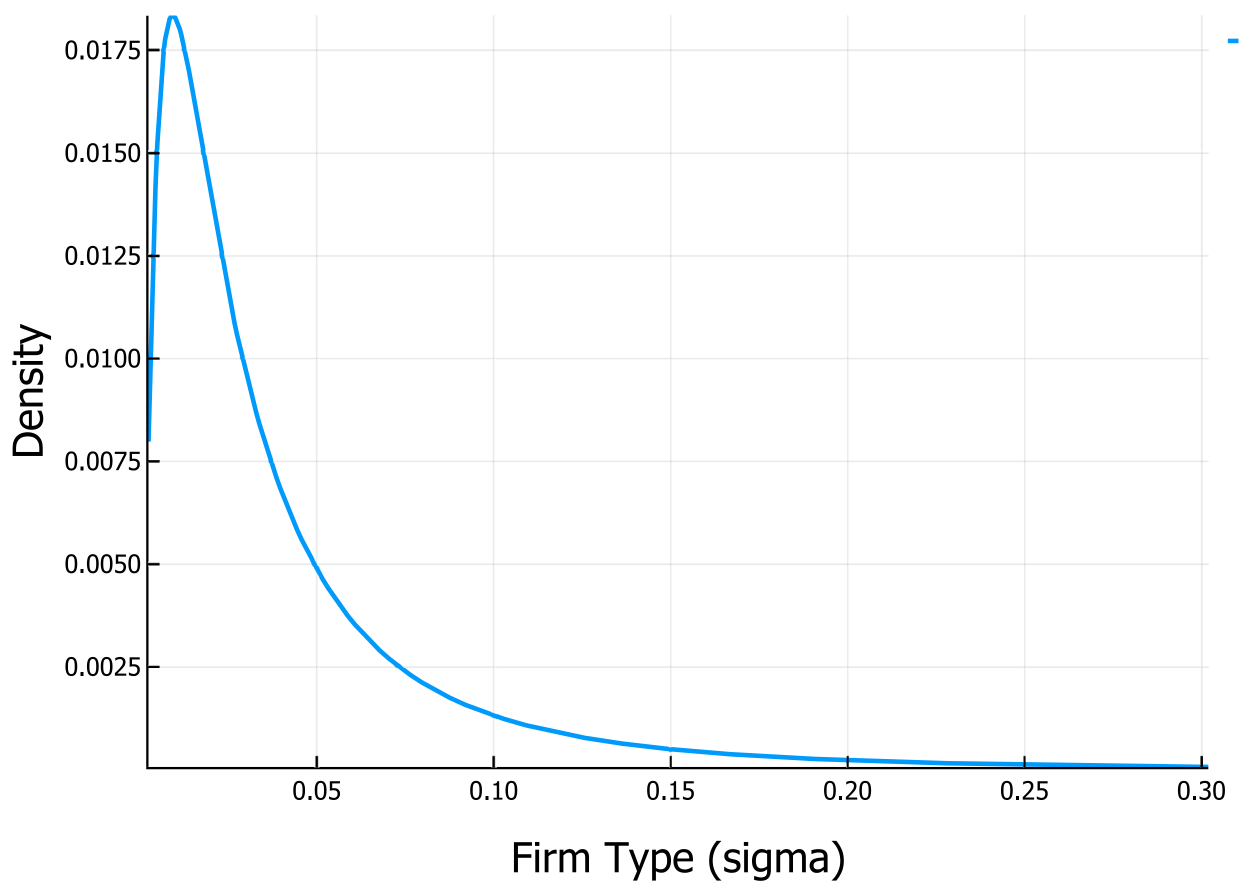

Figure 11: A plot of estimated firm heterogeneity in the variance of their productivity shock. In estimation, the density is discretized between the tenth and ninetieth deciles.

change versus the variable sectors one hundred and fifteen dollar. These costs reflect the profitability of firms having abstracted away from all fixed costs of production and any other overhead not captured in the model.

Regarding model fit, Figure 12 plots the observed and simulated moments while Table 8 provides all seventy-six point estimates. The concise summary is that while the model can do a reasonably good job of matching the observed trends in hours it cannot explain the large variation in observed compensation. The model matches both the level and second moments of hours well excepting the mass point at forty hours per week in the fixed sector. The aggregate search moments and the ratio of fixed to variable contracts can be fit relatively well by the parameters of search technology and the magnitudes of the search costs are largely driven by the implied profits of labor contracts. One explanation of the failure to capture the extremes of the cross-sectional profile in wages is that, conditional on preferences, all workers are equally productive. If workers also had an observable latent type that related to a fixed worker-specific labor productivity then the model might generate additional variation in wages. While a common strategy in the literature is to

Figure 12: Plots Indicating Moment Fit

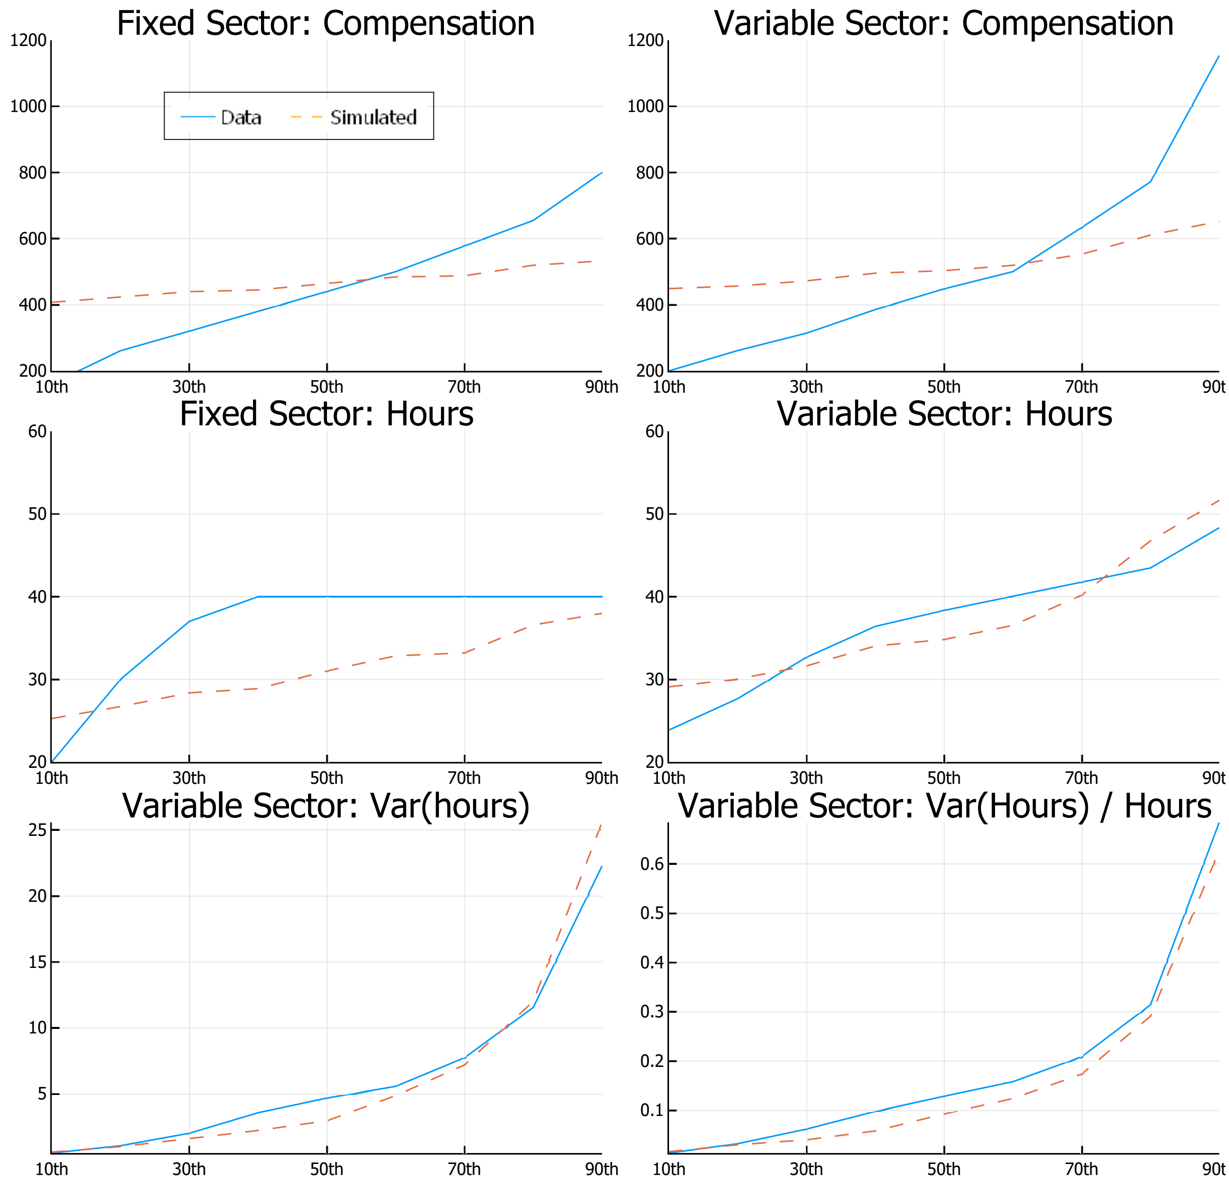

Figure 13: Fit of moments. The blue solid line is the data and the orange dashed line traces out the simulated counterpart.

use a worker's transition between jobs to estimate this fixed productivity this paper avoids using panel data due to the issue of modeling (non-random) transitions in and out of the retail sector. One potentially feasible improvement would be to allow a positive bargaining weight to workers which ties wages to the heterogeneity already existent in the match surplus across worker-firm pairs. This approach is left for future work.

The estimated model fits the higher order moments relatively well. For example, the cross-sectional variance in weekly hours directly relates to firm types to observed labor contracts and suggests that the variability of firm productivity shocks provides a reasonable explanation for observed variation. However, other moments constructed as the ratio of wages to hours and ratio of mean hours to variance in hours pin down the correlation between jobs and avoid an equilibrium where, for example, wages were decreasing in mean hours. I interpret the fit of these moments as evidence that fundamental dynamics of the model are plausible even if, perhaps, the current parameterization does not provide sufficient variation to fit the full cross-section of wages.

VIII. Counterfactual Policy Experiment

One important distinction concerning the scope for welfare improving policy is that the employment contracts in this model are not the best possible among the set of all conceivable contracts. The model has two inherent frictions. The first is a search friction that prevents full employment and sorting. The second and primary friction is that the worker and firm must negotiate a fixed wage before seeing the productivity shocks. As shown in Cooper (1985), an alternative optimal contract contains a shock-contingent wage and hours where the contractual hours and compensation in any period just satisfies the worker's participation constraint under the same bargaining environment found here. This contract allows the risk-neutral firm to write contracts that achieve an efficient level of production while fully insuring the risk-averse worker *ex post*. These contracts involve hours and an effective wage rate that vary with the shock such that the worker receives the same level of utility each period independent of the shock. This contrasts to my environment where I require the worker and firm to fix a wage rate independent of the productivity shock. In my model, the equilibrium wage in the variable sector includes an endogenously determined compensating differential in light of the *ex post* variation in outcomes. That would be absent if wages were allowed to vary. The majority of my sample with hours that vary reports being "an hourly employee" with a fixed wage which implies that they are either risk neutral or that the wage-varying contract is not available. Thus, I make the modeling decision that seems most plausible given the data, but moreover, leaves room for welfare increasing policy that limits the firm's discretion.

Having estimated the model primitives that drive the dynamics of the equilibrium labor market, I now consider the implications of policy designed to reduce the variability workers face in consumption and leisure. Specifically, I conduct a policy experiment that prohibits firms from varying an employee's hours more than eight hours away from that worker's mean hours. The policy entails a mechanical limit on the variation a worker faces and would bind for around twenty-two percent of workers in the sample of observed jobs. This restriction constrains the ability of firms to adjust to productivity shocks and has implications for firm entry, hiring behavior, wages, hours, and employment levels. Accordingly, to evaluate whether the policy improves worker welfare involves determining how the policy changes equilibrium contracts. This determination must account not only for potentially different bundles of hours and consumption but also different job-finding probabilities. The non-linearity implied by changes in firm entry and worker sorting patterns between the sectors and resultant adjustments in wage and hours determination makes an analytical approach to determining the net effects intractable.

My approach to evaluating the policy involves solving and simulating the counterfactual equilibrium and comparing outcomes to the original equilibrium. The policy limits the range in hours which in turn is factored into the expected utility and expected profit calculations of workers and firms, respectively. Workers potentially benefit from reduced variation but this fact also changes the equilibrium wages where the mechanical reduction in variation reduces the magnitude of the compensating differential. Importantly, the latter does not imply altogether lower wages or consumption as changes in mean hours may result in overall higher levels of wages and consumption. The policy constrains the firm's ability to adjust to shocks and to a certain degree moves contracts in the variable sector closer to the environment in the fixed sector. Since firms are solving a constrained version of their original problem, the only way for them to end up better off is for the counterfactual equilibrium outcomes like worker contact rates and sorting patterns to change in their favor. Having solved for the original and counterfactual distributions of matched pairs and unmatched workers and firms, I draw a large sample from both equilibria and compare outcomes for an identical sequence of shocks.

The metric I use to evaluate welfare is the amount of annualized consumption I would need to give the worker in the counterfactual equilibrium to make them indifferent, utility-wise, to the original equilibrium. A positive transfer indicates they are worse off while a negative transfer denotes an improved level of utility. To be clear, the measured transfers are not feasible in the sense that I cannot achieve the counterfactual equilibrium through lump sum transfers as is traditional when, for example, utility is quasi-linear in consumption. The measure of welfare will be an average change in consumption across all workers of the same type which accounts for changes in the probability of employment and in the probability of pairing certain firms.

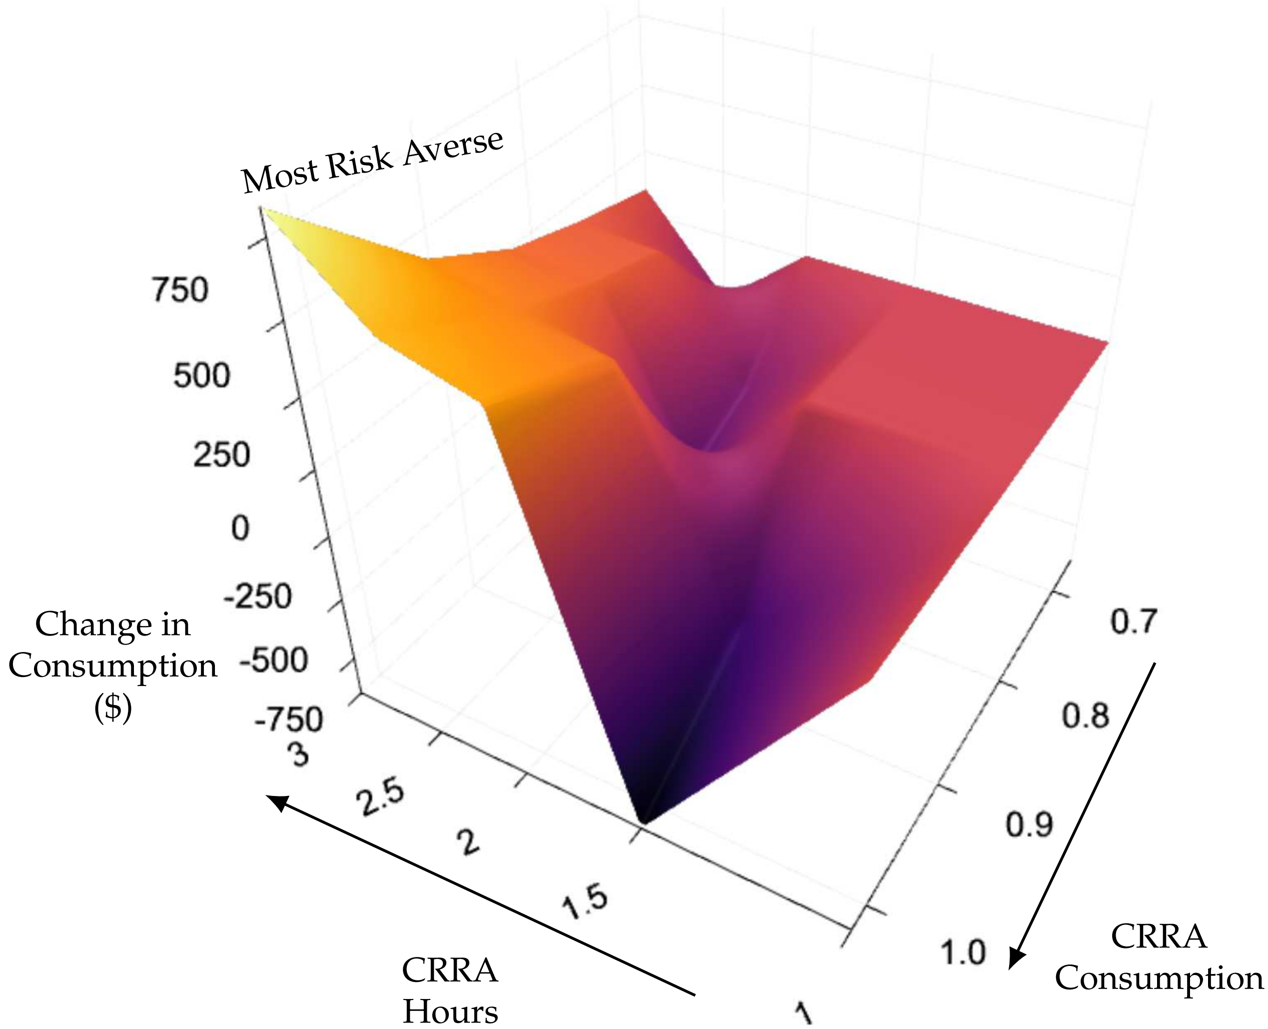

Figure 14: A plot of the change in worker welfare under a policy where firms are prohibited from setting hours that differ by more eight hours from the worker's mean hours. Welfare is measured in the annualized consumption required for each worker type to be indifferent between the original and counterfactual equilibrium. Positive values indicate workers are worse off and negative values indicate an improved average level of utility. For reference, the right hand side plateau sits at approximately zero.

The result of the policy experiment is that the policy binds for around eighteen percent of workers in the counterfactual equilibrium. Aggregate production falls by about two percent which reflects the relatively small but still meaningful constraint on firm behavior. The policy increases the relative proportion of jobs where hours vary as more worker direct their search to the variable sector. Consequently, this movement results in fewer workers participating in the fixed sector. The change in welfare is plotted in Figure 14 as an annualized change in consumption. The average change across all worker types amounts to a forty dollar decrease in annual consumption. This negligible average obscures the large variation where some workers end up eight hundred dollars worse off while others end up as much as nine hundred dollars better off. To contextualize these values, the average pre-tax weekly income of a retail worker in the sample is four hundred dollars meaning the variation is equivalent to two weeks of pay better or worse off under the policy.

Referring to Figure 14, the variation in the change in welfare can largely be understood as a result of changes in the sorting patterns of workers. The workers who end up worse off are the most risk averse workers in the economy. These workers find jobs in the fixed sector under both the original and counterfactual equilibrium. The least risk averse workers are effectively indifferent between the two policies as they start in the variable sector and end up with largely the same outcomes and in the same sector under the policy. In contrast, the workers who see large improvements occupy the middle values of the risk aversion coefficients. They begin in the fixed sector but end up searching and working in the variable sector in the new equilibrium. These workers were presumably not profitable enough to match with high variance firms in the variable sector and so preferred to work in the fixed sector where they had a higher probability of obtaining a job. The policy effectively shields the amount of variation they face, even when matched with a high variance firm, and so they are able to move to the variable sector where a higher vacancy contact rate and higher average contracted utility improves their average welfare. In summary, the policy makes the most risk averse workers worse off largely because through sorting they were not in jobs where lots of variability occurred and end up worse off through the reduced profitability of the industry as a whole. However, the policy does create some winners and while I find the proportion of jobs where hours vary increases, the average amount of variability is brought down by the policy.

IX. Conclusion

This paper formulates and estimates an equilibrium search model to evaluate how limits to employer discretion over hours in the U.S. retail sector impact market efficiency and welfare. To match observed labor contracts and provide an environment where policy intervention may improve outcomes I develop a model of labor contracting where the wage is set before the marginal productivity is known. Firms set hours given the shock and wage which generates uninsured variability in the worker's leisure and consumption across periods. The model specifies firm and worker primitives and allows for bargaining with bilateral heterogeneity as a way to price how control over hours matters in employment contracts within an equilibrium market setting. The search model has two sectors within the labor market that allow for endogenous directed search towards contracts with more and less variability in outcomes in a tractable manner.

To identify the primitives of the model I use a novel two stage approach that allows identification and estimation of worker and firm primitives without data from firms. My strategy involves combining stated preference data collected from workers with data from equilibrium jobs and employment measures specific to the retail sector. In the first stage I estimate worker preferences using data from a purpose-made survey instrument. In the second stage I take these primitives as inputs when estimating the parameters underlying firm and search technology. The estimated model broadly matches observed patterns in wages, hours, and employment in the U.S. retail industry. Using the recovered structural parameters I evaluate the impact of a policy experiment that constrains the extent of variability in hours employers can require from their employees and solve for a counterfactual equilibrium under the new policy. My results indicate that while such a policy appears to reduce variability and have little impact worker welfare, it causes a significant decline in aggregate production and makes the most risk averse workers worse off in equilibrium.

Looking forward, the model appears to capture patterns in hours and higher order moments of the labor contracts but does not fully fit the greater variation found in the cross-section of wages. The issues in fit stem at least partially from parameterization choices and the homogeneity of worker productivity once you condition on preferences. A possible new direction gives workers a positive bargaining weight which would provide them with a share of the match surplus and provide additional variation in wages. This approach may also to improve model fit and inform counterfactual welfare analysis. More generally, further work is needed on how attributes like advanced notice of schedule and flexibility in the timing of hours can mitigate the welfare costs of these types of schedules at a potentially smaller cost to market efficiency. Given the complexity of how such attributes affect the profit of firms, one promising avenue is to use the independent variation generated by randomized control trials, as in Lambert et al. (2015), in conjunction with a structural model of contracting.

References

Barlevy, Gadi and Derek Neal, "Allocating Effort and Talent in Professional Labor Markets," NBER Working Papers 23824, National Bureau of Economic Research, Inc September 2017.

Blau, David M, "Search for Nonwage Job Characteristics: A Test of the Reservation Wage Hypothesis," *Journal of Labor Economics*, April 1991, *9* (2), 186–205.

Bloemen, Hans G., "A model of labour supply with job offer restrictions," *Labour Economics*, 2000, *7* (3), 297–312.

__ , "Job Search, Hours Restrictions, and Desired Hours of Work," *Journal of Labor Economics*, 2008, *26* (1), 137–179.

Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying, "Does Working from Home Work? Evidence from a Chinese Experiment," *The Quarterly Journal of Economics*, 2015, *130* (1), 165–218.

Brown, James N., "Structural Estimation in Implicit Markets," in "The Measurement of Labor Cost" NBER Chapters, National Bureau of Economic Research, Inc, September 1983, pp. 123–152.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin, "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, 2006, *74* (2), 323–364.

Cohen, Alma and Liran Einav, "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, June 2007, *97* (3), 745–788.

Cooper, Russell, "Worker Asymmetric Information and Employment Distortions," *Journal of Labor Economics*, April 1985, *3* (2), pp. 188–208.

__ , John Haltiwanger, and Jonathan L. Willis, "Search frictions: Matching aggregate and establishment observations," *Journal of Monetary Economics*, 2007, *54*, 56–78.

__ , __ , and __ , "Dynamics of labor demand: Evidence from plant-level observations and aggregate implications," *Research in Economics*, 2015, *69* (1), 37 – 50.

Cunha, Flávio and Nick Frazier, "A Model of Human Capital Formation and Contractual Unpredictability," Mimeo, Rice University 2016.

Dey, Matthew S. and Christopher J. Flinn, "An Equilibrium Model of Health Insurance Provision and Wage Determination," *Econometrica*, 2005, *73* (2), 571–627.

Diamond, Peter A. and Jerry A. Hausman, "Contingent Valuation: Is Some Number Better than No Number?," *Journal of Economic Perspectives*, December 1994, *8* (4), 45–64.

Dickens, William T. and Shelly J. Lundberg, "Hours Restrictions and Labor Supply," *International Economic Review*, 1993, *34* (1), 169–192.

Eckstein, Zvi and Kenneth I. Wolpin, "Estimating a Market Equilibrium Search Model from Panel Data on Individuals," *Econometrica*, 1990, *58* (4), 783–808.

__ and __ , "Duration to First Job and the Return to Schooling: Estimates from a Search-Matching Model," *Review of Economic Studies*, 1995, *62* (2), 263–286.

Eriksson, Tor and Nicolai Kristensen, "Wages or Fringes? Some Evidence on Trade-Offs and Sorting," *Journal of Labor Economics*, 2014, *32* (4), pp. 899–928.

Flinn, C. and J. Heckman, "New Methods for Analyzing Structural Models of Labor Force Dynamics," *Journal of Econometrics*, 1982, *18* (1), 115–168.

Flinn, Christopher J., "Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates," *Econometrica*, 2006, *74* (4), 1013–1062.

Golden, Lonnie, "Irregular Work Scheduling and Its Consequences," *Economic Policy Institute*, 2015, *Briefing Paper No. 394.*

—, **Julia R. Henly, and Susan Lambert**, "Work Schedule Flexibility: A Contributor to Happiness?," *Journal of Social Research and Policy*, 2013, *4* (2).

Hanley, Nick, Robert Wright, and Vic Adamowicz, "Using Choice Experiments to Value the Environment," *Environmental & Resource Economics*, April 1998, *11* (3), 413–428.

Heckman, James J., Rosa L. Matzkin, and Lars Nesheim, "Nonparametric Identification and Estimation of Nonadditive Hedonic Models," *Econometrica*, 09 2010, *78* (5), 1569–1591.

Henly, Julia R. and Susan J. Lambert, "Unpredictable Work Timing in Retail Jobs: Implications for Employee Work-Life Outcomes," *Industrial and Labor Relations Review*, 2014, *67* (3), 986–1016.

Ingre, Michael, Torbjörn Åkerstedt, Mirjam Ekstedt, and Göran Kecklund, "Periodic self-rostering in shift work: correspondence between objective work hours, work hour preferences (personal fit), and work schedule satisfaction," *Scandinavian Journal of Work, Environment & Health*, 2012, *38* (4), 327–336.

Joubert, Clement and Petra E. Todd, "Expanding Pension Coverage: A Dynamic Analysis of Chile's Pension Reform," Working Paper 2016.

Keane, Michael and Richard Rogerson, "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom," *Journal of Economic Literature*, June 2012, *50* (2), 464–76.

Lambert, Susan J., "Passing the Buck: Labor Flexibility Practices that Transfer Risk onto Hourly Workers," *Human Relations*, 2008, *61* (9), pp. 1203–27.

—, **Anna Haley-Lock, and Julia R. Henly**, "Schedule Flexibility in Hourly Jobs: Unanticipated Consequences and Promising Directions," *Community, Work and Family*, 2012, *15* (3), pp. 239–315.

—, **Peter J. Fugiel, and Julia R. Henly**, "Precarious Work Schedules among Early-Career Employees in the US: A National Snapshot ," Technical Report, Employment Instability, Family Well-being, and Social Policy Network at the University of Chicago 2014.

—, —, **and** —, "Work Scheduling Study: A Profile Of Retail Sales Associates in a Women's Apparel Firm," Technical Report, The School of Social Service Administration at the University of Chicago September 2015.

Lise, Jeremy, Costas Meghir, and Jean-Marc Robin, "Matching, Sorting and Wages," *Review of Economic Dynamics*, January 2016, *19*, 63–87.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," NBER Working Papers 22708, National Bureau of Economic Research, Inc Sep 2016.

Metrick, Andrew, "A Natural Experiment in "Jeopardy!"," *The American Economic Review*, 1995, *85* (1), 240–253.

Mortensen, Dale and Christopher Pissarides, "New developments in models of search in the labor market," in O. Ashenfelter and R. Layard, eds., *Handbook of Labor Economics*, Vol. 3B of *Handbook of Labor Economics*, Elsevier, 1999, chapter 39, pp. 2567–2627.

Nassauer, Sarah, "Wal-Mart Rolls Out a New Worker Scheduling System," *Wall Street Journal*, August 2016.

Navarro, Salvador and Jin Zhou, "Identifying Agent's Information Sets: an Application to a Lifecycle Model of Schooling, Consumption, and Labor Supply," *Review of Economic Dynamics*, April 2017, *25.*

Postel-Vinay, Fabien and Jean-Marc Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 2002, *70* (6), pp. 2295–2350.

Reynolds, Jeremy, "You Can't Always Get the Hours You Want: Mismatches between Actual and Preferred Work Hours in the U.S.," *Social Forces*, 2003, *81* (4), pp. 1171–1199.

_ **and Lydia Aletraris**, "Pursuing Preferences: The Creation and Resolution of Work Hour Mismatches," *American Sociological Review*, 2006, *71* (4), pp. 618–638.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, January 1982, *50* (1), 97–109.

Stewart, Mark B. and Joanna K. Swaffield, "Constraints on the Desired Hours of Work of British Men," *The Economic Journal*, 1997, *107* (441), pp. 520–535.

Tabuchi, Hiroko, "Abercrombie & Fitch to End On-Call Shifts for Workers," *New York Times*, August 2015.

Thimme, Julian, "Intertemporal Substitution in Consumption: A Literature Review," *Journal of Economic Surveys*, 2017, *31* (1), 226–257.

Tummers, Martijn P. and Isolde Woittiez, "A Simultaneous Wage and Labor Supply Model with Hours Restrictions," *The Journal of Human Resources*, 1991, *26* (3), 393–423.

van der Klaauw, Wilbert, "On the Use of Expectations Data in Estimating Structural Dynamic Choice Models," *Journal of Labor Economics*, 2012, *30* (3), 521 – 554.

van Soest, Arthur, Isolde Woittiez, and Arie Kapteyn, "Labor Supply, Income Taxes, and Hours Restrictions in the Netherlands," *Journal of Human Resources*, 1990, *25* (3), 517–558.

White, Gillian B., "The Very Real Hardship of Unpredictable Work Schedules," *The Atlantic*, April 2015.

A. Appendix: Optimal Labor Contracts

For reference I provide various moments of labor contracts implied by the parameterization in the equilibrium model.

$$
\begin{aligned}
L_v &= \left(\frac{\lambda\alpha e^{\varepsilon}}{w_v}\right)^{\frac{1}{1-\alpha}} \\
\mathbb{E}[L_v] &= \left(\frac{\lambda\alpha}{w_v}\right)^{\frac{1}{1-\alpha}} \exp\left\{\tfrac{1}{2}\sigma^2\left(\frac{1}{(1-\alpha)^2}-\frac{1}{1-\alpha}\right)\right\} \\
&= \left(\frac{\lambda\alpha}{w_v}\right)^{\frac{1}{1-\alpha}} \exp\left\{\frac{\alpha}{2(1-\alpha)^2}\sigma^2\right\} \\
\mathrm{Var}(L_v) &= \left(\frac{\lambda\alpha}{w_v}\right)^{\frac{2}{1-\alpha}} \left(\exp\left\{\frac{(1+\alpha)\sigma_f^2}{(1-\alpha)^2}\right\}-\exp\left\{\frac{\alpha\sigma_f^2}{(1-\alpha)^2}\right\}\right) \\
\mathbb{E}[L_v^{\delta}] &= \left(\frac{\lambda\alpha}{w_v}\right)^{\frac{\delta}{1-\alpha}} \exp\left\{\tfrac{1}{2}\sigma^2\left(\frac{\delta^2}{(1-\alpha)^2}-\frac{\delta}{1-\alpha}\right)\right\} \\
\ln L_v &= \frac{1}{1-\alpha}\left(\ln\lambda+\ln\alpha+\varepsilon-\ln w_v\right) \\
\mathbb{E}[\ln L] &= \frac{1}{1-\alpha}\left(\ln\lambda+\ln\alpha-\frac{1}{2}\sigma_f^2-\ln w_v\right) \\
\mathrm{Var}(\ln L) &= \left(\frac{1}{1-\alpha}\right)^2\sigma_f^2
\end{aligned}
$$

B. Appendix: Additional Information about the UWSF Vigenttes

From each respondent we collect, if employed, their current job characteristics including wage, range of hours, and degree of advanced notice and flexibility. The latter permits verification that the jobs in our vignettes resemble the jobs these individuals hold. We also ask respondents to suppose hypothetically that if in the next month they were to receive thirty jobs offers from employers who hold a vacancy that matches the respondent's qualifications. From that distribution, we ask respondents to report the minimum w^{p}, maximum w^{o}, and most likely wage w^{r} from which we can construct a triangular distribution of expected job offers. These moments are then used to adapt the wage offers from the jobs into offers that necessarily correspond to an individual-specific set of of low w_{lo}, low to mid w_{m1}, mid to high w_{m2} and high w_{hi} wage offers. Each job has one of four possible wage levels created from the following set of equations:

$$w_{i,lo} = 0.9 \times w_{i,p} + 0.1 \times w_{i,r}$$

$$w_{i,m1} = 0.6 \times w_{i,p} + 0.4 \times w_{i,r}$$

$$w_{i,m2} = 0.4 \times w_{i,r} + 0.6 \times w_{i,o}$$

$$w_{i,hi} = 0.1 \times w_{i,r} + 0.9 \times w_{i,o}$$

C. Appendix: Algorithm for Finding Steady State Equilibrium

First I discretize the set of κ and σ into N_κ and N_σ intervals between the tenth and ninetieth quantiles. I limit myself to the inner deciles to avoid letting types in the tail drive dynamics as behavior in the tails is primarily identified through parametric assumption.

1. Calculate θ_m and $\eta_m(\kappa)$ for both markets assuming all workers participate in that market. Label the market with the higher θ_m as Market A and the other as Market B. Let $\mathscr{A}$ ($\mathscr{B}$) be the set of workers in Market A (B).

    - In practice, this will almost always be the variable market as a match for any given κ and σ always has a higher expected profit in the variable market. Intuitively, a firm in the fixed market is solving a constrained version of the problem in the variable market.

2. Take the least profitable worker, κ_j, in $\mathscr{A}$ and compute θ_B and $\eta_B(\kappa_j)$ for Market B where only κ_j participates.
3. Check if $U^B(\kappa_j) \geq U^A(\kappa_j)$ using the θ_m and $\eta_m(\kappa_j)$ where j participates in both markets.

- Note that if any firm accepts j in the fixed market then all will since R^f and Q^f are independent of firm type.
- Without loss of generality, to break ties I assume a worker prefers the variable market if and only if $U^v(\kappa_j) > U^f(\kappa_j)$.

4. If so, put $\kappa_j \in \mathscr{B}$ and take her out of $\mathscr{A}$. Then find the next least profitable worker in $\mathscr{A}$ and repeat from step 2 until $U^B(\kappa_j) < U^A(\kappa_j)$

   - As workers transfer from $\mathscr{A}$ to $\mathscr{B}$ both θ_A and θ_B will rise. Market A has higher expected profit as the least profitable κ has left. Market B is also has higher expected profit as the new κ will always be more profitable than the other workers already in $\mathscr{B}$. More profitable markets encourages firm entry which drives the vacancy contact rates and thus the value of unemployment.

The resultant $\theta_v, \theta_f, \eta_v(\kappa)$, and $\eta_(\kappa)$ imply a division of workers between the two markets where optimizing behavior in directed search by workers implies a worker cannot be better off by searching in the other market. Using this division and θ_m and $\eta_m(\kappa)$ we know the steady state distribution of unmatched workers F_κ^m for both markets. Integrating against the discretized worker types we know u_m which implies the measures of vacancies $v_m = \theta_m u_m$ and the distribution of matches. To simulate the sample of observed jobs I draw a large sample from the distribution of equilibrium matches and calculate the wage, hour, and worker flow moments.

D. Appendix: Additional Tables and Figures

Table 10: Types of Retailers in Sample of Retail Workers

%	Description	Census ID
10%	Autombile Related	4670–4690
17%	Furniture, Appliances, & Hardware	4770–4890
30%	Food and Drug	4970–5080
6%	Gas Stations	5090
13%	Clothing	5170–5280
16%	Music, Book, & Department stores	5290–5390
8%	Other	5470–5580

A further breakdown of the 2002 Census Industry codes for the 405 retail jobs in the sample.

Table 11: Reported Job Attributes in the UWSF Survey

Flexibility	None	A little	Some	A lot
All	47%	24%	21%	8%
Hourly	56%	26%	14%	3%
Female	48%	23%	22%	8%
Male	47%	26%	19%	8%
Part-Time	48%	26%	17%	9%
Full-Time	47%	23%	23%	7%
Advanced Notice	$\leq$ 1 Week	1–2 Weeks	3–4 Weeks	4+ Weeks
All	37%	21%	6%	36%
Hourly	38%	26%	6%	30%
Female	37%	22%	6%	35%
Male	36%	20%	6%	38%
Part-Time	45%	32%	7%	17%
Full-Time	32%	16%	6%	46%
Variability in Hours	None	$\pm$10%	$\pm$50%	Mean IR
All	30%	70%	35%	0.7
Hourly	31%	69%	37%	0.73
Female	30%	70%	36%	0.75
Male	31%	69%	34%	0.63
Part-Time	16%	84%	60%	1.21
Full-Time	38%	62%	23%	0.45

Attributes of reported jobs from Understanding Work Schedule Forecastability sample. This table suggests that the attributes contained in our hypothetical jobs resemble those of equilibrium jobs. *IR*, or instability ratio, is a measure of the degree of variation in hours for any particular job defined as the range of hours divided by usual hours.

Table 12: Additional UWSF Sample Statistics

	Females	Males	Part-Time	Full-Time	All
General					
% of Total Sample	40%	32%	24%	48%	72%
Male	0%	100%	47%	43%	44%
Mean Age	32.9	33.1	32.9	33.0	33.0
Cohabitating	64%	68%	60%	68%	66%
Currently a Student	9%	10%	13%	8%	9%
Any Children	53%	63%	53%	59%	57%
Any Children Under 6	47%	37%	47%	41%	43%
Total Time on Survey	31.7	29.9	27.1	32.8	30.9
Employment					
Number of Jobs	1.5	1.4	1.7	1.4	1.5
Mean Usual Hours	36.2	33.4	20.3	42.3	35.0
Mean Min Hours	29.5	26.3	14.7	34.7	28.1
Mean Max Hours	50.4	43.3	34.3	53.7	47.2
Suppose you recieved 20 job offers in the next 2 months, what do you think would be the…					
Lowest Expected Wage	17.6	14.2	12.9	17.7	16.1
Most Likely Wage	28.3	20.1	19.9	27.0	24.6
Highest Expected Wage	40.4	29.1	29.8	38.2	35.4
Since you turned 20, how many years were you primarily a…					
Full-time Worker	8.5	7.9	6.5	9.1	8.2
Part-time Worker	1.7	1.7	3.2	0.9	1.7
Student	1.3	1.9	1.7	1.6	1.6
Other	0.7	0.7	1.2	0.4	0.7
Race and Ethnicity					
Black	8%	11%	11%	9%	10%
Hispanic	9%	9%	9%	9%	9%
White	83%	85%	82%	85%	84%
Highest Degree Attained					
Less Than HS	3%	2%	2%	2%	2%
High School	38%	35%	42%	34%	37%
A.A. Degree	10%	15%	14%	11%	12%
B.A. Degree	32%	32%	28%	34%	32%
Graduate Degree	17%	17%	14%	19%	17%
Sample Size	413	329	247	495	742

Additional sample statistics for the UWSF sample. Part-time is defined as reporting fewer than 35 usual hours per week in their primary job.

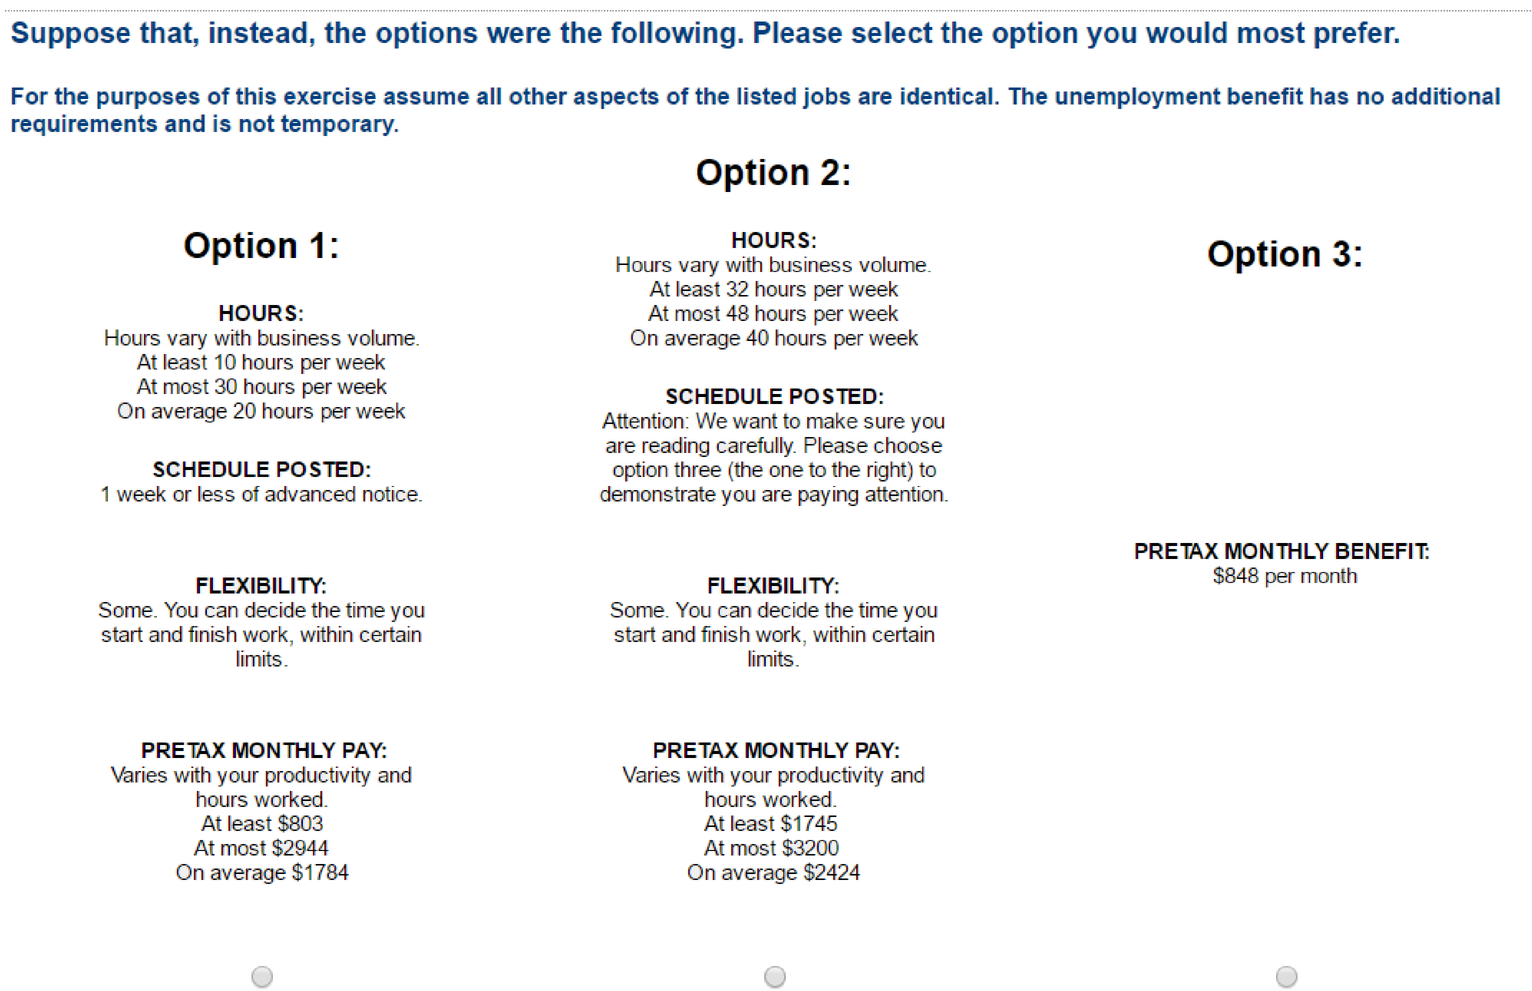

Figure 15: The Attention Filter from the Understanding Work Schedule Forecastability survey. Every survey-taker saw the above as the 12th of 14 vignettes. In Option 2 the description for "Schedule Posted" instructs them to disregard the above instructions and choose Option 3. For reference, across vignettes the unemployment benefit is empirically the least likely option to be chosen. Filtering on whether they answered this question correctly leaves 40.10 of original respondents.

Table 13: Unweighted Estimated Preference Parameters

Parameterization

$$\frac{C^{1-\rho_{c,i}}-1}{1-\rho_{c,i}} - \delta\frac{L^{1-\rho_{l,i}}-1}{1-\rho_{l,i}} + \psi X - \mathbb{1}_{\text{ue}}\overline{g} \qquad \begin{pmatrix} \ln\rho_c \\ \ln\rho_\ell \end{pmatrix} \sim N\left(\begin{pmatrix} \rho_{c0} \\ \rho_{\ell 0} \end{pmatrix}, \begin{bmatrix} \rho_{c1}^2 & \cdot \\ \rho_{c\ell}\rho_{c1}\rho_{\ell 1} & \rho_{\ell 1}^2 \end{bmatrix} \right)$$

Re-weighted Estimates

	δ	ρ_{c0}	$\rho_{\ell 0}$	ρ_{c1}	$\rho_{\ell 1}$	$\rho_{c\ell}$	$\overline{g}$
Estimate	0.65	-0.20	0.56	0.21	0.50	0.58	0.79
SE	(0.19)	(0.01)	(0.14)	(0.01)	(0.09)	(0.01)	(0.16)

	Advanced Notice		*Flexibility*		
ψ :	1–2 Weeks	3+ weeks	A little	Some	A lot
Estimate	0.26	0.31	0.54	0.82	1.23
SE	(0.04)	(0.04)	(0.05)	(0.05)	(0.04)

N=1026 Log-Likelihood: -927.7

Table 14: Results for preference estimation without the re-weighting to match the NLSY97 sample.

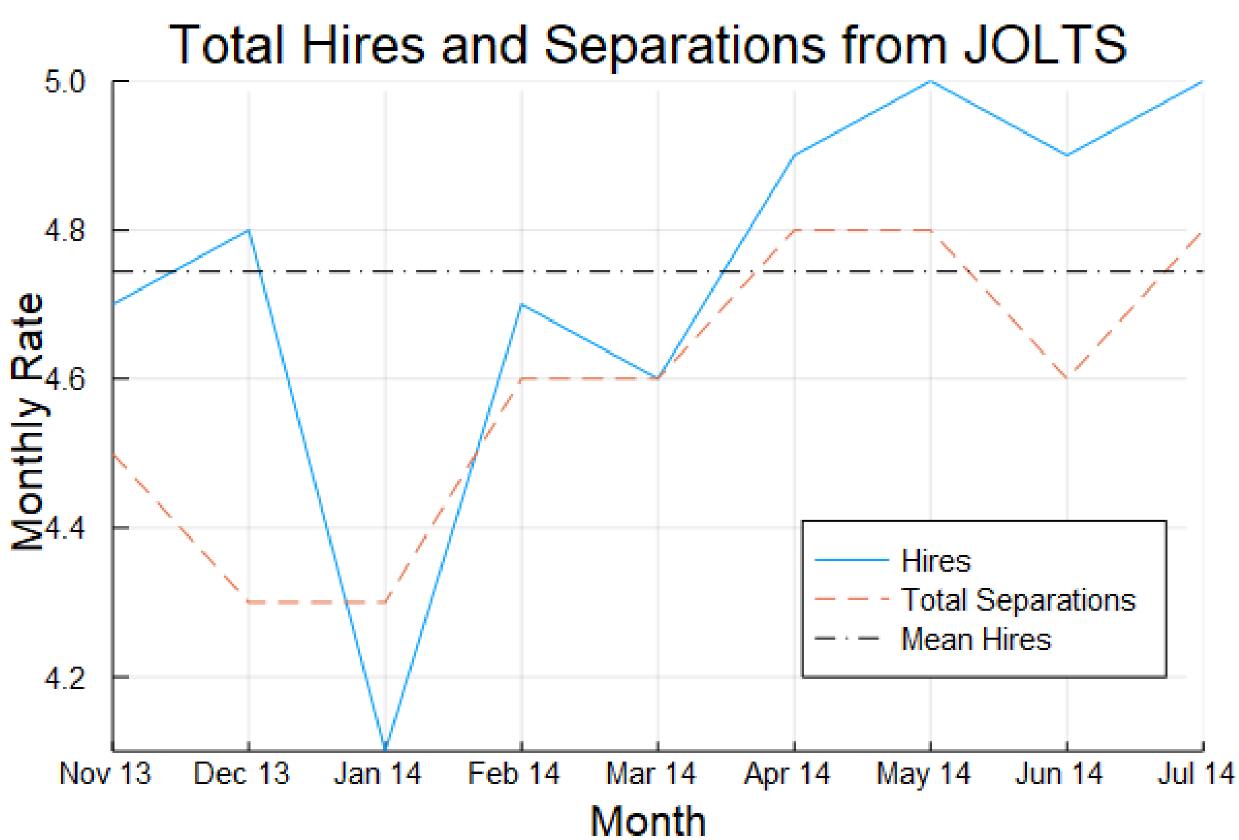

Figure 16: The Hires and Total Separations series for November 2013 through July 2014 from the JOLTS database. Rates are seasonally-adjusted.

Table 15: Comparison of Simulated and Observed Moments

Aggregate Moments									
	UE Rate	θ	Job Openings	$\frac{\text{variable}}{\text{fixed}}$					
Data	0.06	0.42	0.03	1.18					
Simulated	0.09	0.41	0.08	1.01					
Fixed: Wages									
Data	7.68	8.25	9.00	10.00	11.07	13.00	14.72	16.50	20.00
Simulated	10.72	10.77	10.84	10.91	11.01	11.27	11.48	11.76	12.17
Fixed: Hours									
Data	20.00	30.00	37.00	40.00	40.00	40.00	40.00	40.00	40.00
Simulated	26.39	29.22	31.36	33.02	35.43	36.33	37.08	37.72	38.27
Fixed: Wages / Hours									
Data	0.22	0.24	0.28	0.30	0.34	0.38	0.41	0.47	0.76
Simulated	0.28	0.29	0.29	0.30	0.31	0.34	0.37	0.40	0.46
Variable: Wages									
Data	8.00	8.69	10.00	11.00	12.00	14.00	15.47	18.87	24.00
Simulated	10.45	10.48	10.51	10.54	10.57	10.61	10.64	10.67	10.76
Variable: Hours									
Data	23.87	27.67	32.67	36.40	38.33	40.00	41.67	43.47	48.33
Simulated	37.98	38.75	39.16	39.51	39.83	40.22	40.49	40.84	41.22
Variable: Var(Hours)									
Data	0.50	1.06	2.00	3.56	4.67	5.56	7.72	11.53	22.26
Simulated	0.44	0.91	0.99	1.70	2.65	3.96	4.57	9.05	14.89
Variable: Wages / Hours									
Data	0.24	0.27	0.29	0.32	0.35	0.39	0.43	0.48	0.72
Simulated	0.25	0.26	0.26	0.26	0.27	0.27	0.27	0.28	0.28
Variable: Var(Hours) / Hours									
Data	0.01	0.03	0.06	0.10	0.13	0.16	0.21	0.31	0.68
Simulated	0.01	0.02	0.02	0.04	0.07	0.10	0.11	0.24	0.39

The observed and simulated moments outlined in Table 8.

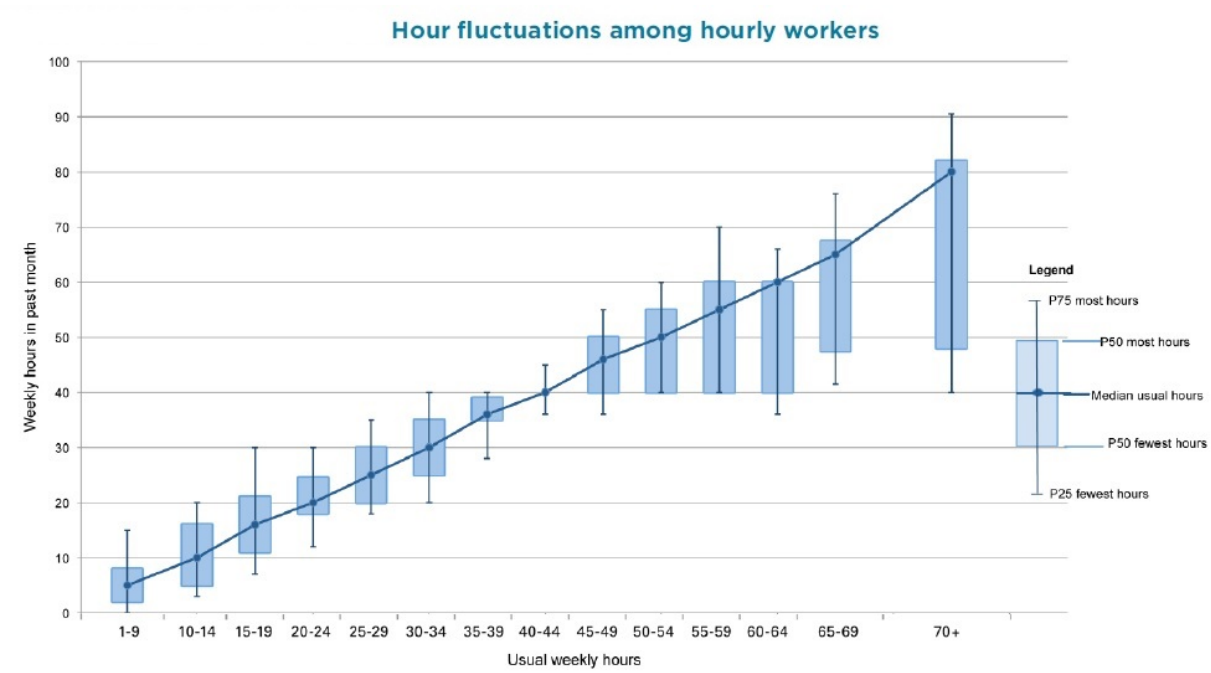

Figure 17: The distribution of actual hours worked in a week conditional on usual hours for employees in their early 30s. Taken from Lambert et al. (2014) who use information from the 2011 round of the National Longitudinal Survey of Youth 1997.