A Model of Human Capital Formation and Contractual Unpredictability

Current Version

Flávio Cunha Nick Frazier

Rice University

February, 2016

We study the effect of restricting the contracting space within a partial equilibrium environment. We develop a model of human capital accumulation where jobs have multiple characteristics including variability in hours and wages. We highlight the difficulties of using standard panel data to estimate and distinguish between worker preferences and the primitives of the offer function — an issue with serious implications when comparing employee welfare. We propose using stated preferences data to trace out worker preferences over a variety of job attributes which we then take to the NLSY97 to recover the parameters underlying the distribution of job offers allowing for dependence on accumulated human capital. This approach provides a richer characterization of the process and permits us to quantify the effect of restricting the offer space on the labor force directly.

As new technology in on-demand scheduling expands the ability of employers to tailor labor hours to day-to-day consumer demand, evaluating changes in employee welfare requires developing models beyond the standard wage/hour paradigm. One often-cited consequence of on-demand scheduling is an increase the variability of earnings, work schedules, and even employment.[1] However, if these conditions make employees significantly worse off then why do such contracts exist between freely contracting workers and firms? Lambert et al. (2014) underscore the issue's salience for the U.S. labor force by using the National Longitudinal Survey of Youth 1997 (NLSY97) to estimate that of the 62% of the employed labor force who are hourly, around 41% receive one week or less of advance notice for their work schedule.[2] To help quantify the effect on employee welfare we look at the effect of restricting the contracting space within a

[1] Two recent surveys of the issue are Lambert et al. (2014) and Golden (2015). Examples of other recent works 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).

[2] For perspective, the percentage for non-hourly employees is 33% and together they represent 38% of the total employed labor force. These numbers are taken from Table 3 of Lambert et al. (2014). Note that respondents were in their early 30's.

partial equilibrium environment.

We develop a model of human capital accumulation where jobs have multiple characteristics and variability in hours and wages. The standard panel data set records accepted jobs, the characteristics of the jobs, and the characteristics of the people who hold them. With sufficient structure on the utility and offer functions, parameters can, in principle, be estimated that rationalize observed distributions for any particular job attribute. Essentially, if we see very few jobs with both high variability in hours and high flexibility then estimation finds a solution that puts low probability on that observation. However, this framework cannot separately identify whether such offers are scarce because people will not accept them or because they are not often offered. We propose using stated preferences data to trace out worker preferences over a variety of job attributes which we then take to the NLSY97 to recover the parameters underlying offer functions that depend on accumulated human capital. This approach provides a richer characterization of the process and permits us to quantify the effect of restricting the offer space on worker welfare directly.

On-demand scheduling may reduce labor costs and provide greater opportunity for employment. However, when describing the efficiency of capital it is worth remembering that employees have preferences over the allocation and variability of their scheduling. By contractual unpredictability, we have in mind jobs with unpredictable scheduling and/or compensation as opposed to being temporary in nature or where the nature of the work to be done changes over time. If individuals prefer predictable schedules then the worker is only indifferent between two jobs with differing levels of predictability when the wage for the less predictable job is higher. The wage difference required for indifference is the "compensating differential" of Rosen (1987). A friction-less equilibrium model implies that observed jobs are those where the contract satisfies the employee's participation constraint. To include the possibility of non-optimality we introduce frictions in our search model. Restricting the contracting space is optimal in the presence of unrealized externalities or if firms have monopsony power (see Manning (1995)). The latter is implied by search frictions and match specific capital both of which are features of our model.

Several papers have studied jobs that consisted of a wage and a set of hours in an equilibrium framework in order to explain observed variation in hours. Recent examples include Blundell et al. (2008) who look at how hours evolve during job changes. Bloemen (2008) who uses a

search model where offers are a wage/hour tuple and uses stated preferences on desired hours to estimate an implied compensating differential. Cooper et al. (2007) develop a theory of firm with many workers and variable hours but the decisions of the worker is largely unmodelled. To our knowledge there is no equilibrium model with variability and elements of unpredictability in scheduling.

The plan of the paper is as follows. In Section 2, we develop a model of human capital accumulation when jobs have multiple attributes. Section 3 discusses identification given data and details our approach. Section 4 describes the vignettes and methodology used to estimate individual preferences. Section 5 contains the econometric framework and approach used to estimate the human capital accumulation model. In Section 6 we presents the results of estimation and in Section 7 we use these estimates to characterize the welfare costs of policies. Section 8 concludes.

II. Model

In our model an individual enters the labor force and accumulates human capital and endogenously transitions in and out of the labor force and from job to job over a finite work life of discrete periods. Every period our agent chooses between keeping her current labor force status, moving to a new job, or leaving the work force. She receives a known and fixed number of offers every period though no *a priori* restriction is placed on the number per period. In a dynamic sense, offers depend on accumulated work experience but are otherwise independent of individual characteristics. Labor contracts will be multi-dimensional but fixed for the duration of the job. The reward for work depends on the consumption value of earnings, the (dis)utility of labor and other job characteristics, and the discounted value of increasing work experience. The final element casts the problem as a dynamic discrete choice model similar Keane and Wolpin (1997) and others.

A. Job Offers

Define an offer as a tuple, (W, H, z) where W and H are a distribution of wage and hours, respectively, and z is a vector of other job characteristics. For $s = \{W, H\}$, define $\mathbb{Q}_s$ as a class of distributions with moments θ_s where one element of θ_s is a finite mean. Then $H \in \mathbb{Q}_H$ and $W \in \mathbb{Q}_W$. Let z be a $K_z \times 1$ vector of binary elements where $z_i = 1$ implies the job has characteristic

i.[3] Finally, define $\mathbb{J}\colon \mathbb{Q}_W \times \mathbb{Q}_H \times \{0,1\}^{K_z}$ and denote a generic offer as $q = (W, H, z)$, an element in $\mathbb{J}$. Given the space of job offers, we introduce $F(\cdot;\theta_{\mathbb{J}}|X)$ as the distribution of job offers such that $F\colon \mathbb{J} \to [0,1]$ given X. The variable X contains characteristics that affect the probability of a certain offer. While in our application we focus on modeling work experience, we allow X to contain time invariant proxies for initial human capital, such as education, early work experience, and other relevant factors. Notably, in the data we observe the distribution of accepted offers, $G(q, X)$. We parameterize F with $\theta_{\mathbb{J}}$ and the agents take $F(\cdot;\theta_{\mathbb{J}}|X)$ as known and exogenous.

We abstract away from negotiation over job contracts by assuming some exogenous division of bargaining power and a complete information environment. Let $\alpha \in (0,1)$ represents the worker's share of the surplus generated by employment. We assume the firm owns some employer-specific and time invariant technology and makes the optimal offer given the employee's unobserved employer-specific productivity. The employer observes the work experience of the prospective employee and can verify effort which results in take-it-or-leave-it offers. This framework models a complete contract over all states of the world where variability in hours and wages is possible but forecastable and specified in the job contract. Firms do not observe nor respond to other offers from other firms. We do not further model the employer's decisions.

B. Description of Workers

From the worker's perspective timing proceeds as follows. She enters the workforce at some age, A_0, as unemployed with some initial human capital, X_0. At the beginning of the first and each period thereafter she receives one or more draws from an offer distribution. She then chooses whether to accept an offer and begin work immediately, transition to home production, or keep her current status for the period. The current status at time t is summarized in a state vector, Ω_t, and contains the job currently held q_t, accumulated work experience X_t, the current period t, and any other exogenous and time invariant factors contained in X_0. The reward in any particular period can flexibly depend on other factors but we assume offers depend solely on the worker's accumulated human capital.

A worker who rejects all offers chooses home production and consumes a monetary benefit not tied to any labor hours or other job characteristics. In the framework of an offer we can

[3] In our application, we will construct z as a concatenation of sub-vectors where one of multiple possible values is selected. For example, a job could have no, a little, some, or a lot of flexibility and the 4×1 sub-vector will always indicate which of the four options describes the job. See Section IV for more information.

characterize home production as the offer $q^u = (W^{UE}, 0, \vec{0})$ where W^{UE} is a degenerate distribution of compensation. This compensation may depend on the state-vector, e.g. accumulated work experience or their current status.[4] One significant issue in a finite period situation is how to model an involuntary transition to home production. Clearly we observe involuntary unemployment in the data but incorporating exogenous termination within discrete time and in the context of the model has an uncertain interpretation. For example, should we equate several brief spells of employment with continuous unemployment? What percentage of observed job to job transitions were involuntary? Given our application and chosen counterfactuals we assume transitions are voluntary but allow home production to depend on work experience.[5]

Every worker has the same endowment of hours to split between labor and leisure. We set the endowment to maximum value of hours from an offer in H and further normalize this value to 1. The worker spends l_t of these hours laboring for the hours specified in the contract and the rest on leisure. For a contract with variable hours we assume their reward for working is the expected utility over the distribution of hours. The analogous assumption is made for the contracts with a distribution of wages. This approach coincides with the employee choosing between the expected utility of alternatives. Therefore, the actual realization of hours and compensation could only affect the updating process for work experience which we discuss below.

Every period our agent receives O new offers and may select one of the O offers, keep her current status and, if employed, transition to home production. We index the finite set of alternatives by $m \in M$ for $M = \{0, 1, \ldots, O+1\}$. Let $p_t^m = 1$ indicate that she opts for choice m in period t and that $p_t^k = 0,\ \forall k \neq m$. For convenience let $p_t^0 = 1$ indicate she prefers home production and $p_t^1 = 1$ that she prefers her current status. Notice that a currently unemployed individual receives the same reward for $p_t^0 = 1$ as for $p_t^1 = 1$. In our dynamic discrete choice framework, p_t^m serves as the policy function of our agent and the solution to her optimization problem. Finally, let $p_t = \{0, 1\}$ summarize whether she works in period t.

The work experience an individual accumulates potentially affects the offers received. The reward for a contract includes the future discounted value of holding that job next period as well as the value of expected offers given the updated human capital. Since hours are heterogeneous across jobs, individuals will be heterogeneous in accumulated experience as well. Furthermore, as

[4] Potentially the value of home production in period t could depend on q_{t-1} but not on $q_{t-s}, s > 1$ but new offers do not.

[5] In our application human capital will not depreciate which, for a long enough work-life and a sufficiently generous $q^u(X_t)$, implies retirement for some workers.

some workers will face a distribution of hours we update their work experience, x_{t+1} by

$$x_{t+1} = \begin{cases} x_t + \overline{h} & \text{if } p_t = 1 \\ x_t & \text{if } p_t = 0 \end{cases} \tag{1}$$

where, as above, $\overline{h}$ is the mean of the hours worked distribution. Using the mean of hours in the law of motion for updating x_{t+1} is consistent with what a worker might report on their resume/CV. Since we work within a finite period framework that will correspond to year in the data, we feel comfortable using $\overline{h}$ to capture the updated experience of the worker.

c. Reward for Work

Let $U(c_t, l_t, z_t)$ be a concave utility function that maps $\mathbb{R}_+ \times \mathbb{R}_+ \times \{0,1\}^{K_z} \to \mathbb{R}$, and which, in our model, will correspond to a surjective function over $\mathbb{J} \to \mathbb{R}$. We impose that $\frac{\partial U}{\partial c_t} \geq 0$ and $\frac{\partial U}{\partial l_t} \leq 0$. All individuals have identical reward functions, including a common discount factor, β, though all inputs will vary by individual. We suppress the i subscript below for ease of notation. Without loss of generality $U(\cdot)$ may also be a function of elements in Ω_t other than X_t. However, since (W, H, z) fully describes a job, any additional characteristics that enter $U(\cdot)$ linearly will prove irrelevant for deciding between alternatives. Therefore, we omit any additional arguments into utility here for ease of notation.

Denote an individual's work life as the set of periods $\{1, \ldots, T\}$. The reward of employment depends on the attributes of the job chosen which may include random outcomes of a known distribution. The individual also forms expectations of the value of possible future offers conditional on whether they worked this period and the current state vector. The reward for home production is not stochastic given accumulated work experience. Since there is no storage technology and therefore no savings, consumption equals compensation every period. The value of home production is then

$$\begin{aligned} V^0(\Omega_t, t) &= U(W^{UE}, 0, 0) + \beta \, \mathbb{E}\left[V^0(\Omega_{t+1}, t+1) \middle| \Omega_t, p_t = 0 \right] \qquad \text{for } t < T \\ V^0(\Omega_T, T) &= U(W^{UE}, 0, 0) \end{aligned} \tag{2}$$

where, here, Ω_{t+1} contains $x_{t+1} = x_t$ and the expectation is over possible future offers. The reward

for any particular job equal the expected value of this period's reward plus the value of holding that job and additional work experience next period. An employed individual faces $O+1$ possible jobs indexed by j and each with a specific reward and discount future value equal to

$$\begin{aligned} V^j(\Omega_t, t) &= \mathbb{E}_{W,H}\left[U(c_t, l_t, z_t) + \beta\, \mathbb{E}\left[V^j(\Omega_{t+1}, t+1)\Big|\, \Omega_t, p_t = 1\right]\right] \qquad \text{for } t < T \\ V^j(\Omega_T, T) &= \mathbb{E}_{W,H}\left[U(c_T, l_T, z_T)\right] \end{aligned} \tag{3}$$

where $c_s = w_s l_s$.

In any period, an employed worker receiving O offers chooses between $O+2$ alternatives contained in M. The agent's value function can be written as the maximum over the alternative-specific value function, or

$$V(\Omega_t, t) = \max_{m \in M} \left\{V^m(\Omega_t, t)\right\}. \tag{4}$$

Each alternative-specific value function has the properties of a Bellman equation. We assume that ties are broken arbitrarily.

D. Evolution of Offer Function

The probability of any particular job depends on an individual's accumulated human capital. To complete our model we need to impose some structure on how the offer distribution, $F(q|X_t)$, depends the current work experience, X_t. If the space of realized wages and hours were discrete and finite, like that of other job characteristics, then we could conceive of a large matrix where for each possible value of X_t we had a vector of probabilities for every possible combination. At this point we allow for a more general framework where the moments of the offer distribution are a direct function of X_t. The collected moments must then be sufficient statistics for the distribution but the functional dependence can be somewhat flexibly specified

Let $\theta_{\mathbb{J}}$, a $K \times 1$ vector, be the moments of F. So far we have imposed little structure on $\theta_{\mathbb{J}}$ other than that the distributions W and H have a finite mean. We want to allow for the case that $F(q|X) \neq F(q|\tilde{X})$ for some $X \neq \tilde{X}$ and we assume this arrives from a functional dependence that

updates $\theta_{\mathbb{J}}$. Define $\theta_{\mathbb{J}|X_t}$ as the expected values of $\theta_{\mathbb{J}}$ given X_t where we assume that

$$\theta_{\mathbb{J}|X_{t+1}} = \Theta(\theta_{\mathbb{J}|X_t}, X_{t+1}, \varepsilon_t; \theta_\Theta). \tag{5}$$

where ε_t is an independent, serially uncorrelated, and mean-zero innovation realized in period t. Thus, the individual updates their expected future offers using the moments of the current offer function $\theta_{\mathbb{J}|X_t}$, a possible future work experience X_{t+1}, and a realized innovation ε_t. Adding the error term increases the flexibility of the updating function but requires that the individual also integrate against the possible future values of ε_{t+1}. The updating function need not necessarily be continuous in any of its arguments.

The above laws of motion and the innovation terms are known by individuals at time t and they further take them as exogenous and given. From our perspective, $\Theta(\cdot)$ and its θ_Θ parameters are unknown and all innovations are unobserved. Though the solution is deterministic from the point of view of the agent, the unobservability of the errors makes the solution to (4) probabilistic from the econometrician's point of view. We discuss the limitations this feature imposes on inference below.

III. Data and Identification

We are interested in identifying the impact of restricting the contracting space on labor outcomes and welfare. We model a partial equilibrium environment of individual behavior where agents optimize taking the characteristics and probability of a particular job as exogenous and immutable conditional on work experience. Knowing the primitives of the policy and value functions allow us to evaluate the cost of restricting the space of contracts where optimization behavior endogenously drives distortions and permits a rigorous interpretation of parameters estimated. Studying equation (4) reveals that knowing the set of alternative-specific values at each period is equivalent to knowing both the policy and value functions.

Panel data typically consists of a history of chosen jobs and the *ex-post* state vector for each respondent. In other words, over the work-life of our agent we will observe the attributes of both

the job and individual for all periods. Keeping this feature in mind, we re-write equation (3) as

$$V^j(\Omega_t, t; \theta) = \mathbb{E}_{W,H}\left\{U(c_t, l_t, z_t; \theta_U) + \beta\, \mathbb{E}\left[V^j(\Omega_{t+1}, t+1)\Big|\, \Omega_t, p_t = 1; \theta_\Theta\right]\right\} \tag{6}$$

which illustrates our three primary objects of interest. The first is θ_U, which parameterizes the utility function. The second is θ_Θ which parameterizes the offer distribution. Given some structure on the expectations, the third is simply the discount factor, β. For notational convenience, define $\theta = \{\theta_U, \theta_\Theta, \beta\}$. Clearly, the state vector over the optimal path of any agent is data and so the issue will be the inference over other alternatives not chosen.

In general, recovering θ from (6) is straightforward given sufficient structure. Let $G(q, x)$ denote the joint distribution of jobs and work experience observed in the data. If $G(\tilde{q}, \tilde{x})$ is small for some $(\tilde{q}, \tilde{x})$ then finding a combination of functional forms for $U(\cdot)$ and $F(\cdot)$ along with a $\hat{\theta}$ to replicate that feature should, in principle, not prove too great a challenge. This argument extends to the universe of observed (q, x). The concern, however, is the implicit inability to separately identify the structure of $U(\cdot)$ from $F(\cdot)$. This issue need not be universal but the interpretation of our counterfactual, exogenously restricting the space of $q \in \mathbb{J}$, hinges on whether certain job attributes contribute significantly to employee preferences or whether, in contrast, their presence or absence is largely a function of how often such an offer occurs. For example, a policy that restricts a disliked but common job attribute could improve employee welfare but were it to restrict a benign and common attribute then in our model the fall in offers potentially reduces employee welfare. Thus separating the contribution of $U(\cdot)$ from $F(\cdot)$ to $G(\cdot)$ proves crucial to our application.

Given our focus, the principle obstacle to identification is the dynamic set-up of our model. We could assume the offer distribution to be invariant over the work-life of the individual at the cost of potentially missing a large part of market dynamics. Generally, $G \neq F$ and the substantive differences could significantly impact the interpretation of our counterfactuals. Observing this motivation we propose the following identification strategy. First, we seek to identify the parameters in $U(\cdot)$ using a survey asking respondents to choose between hypothetical job offers in an environment constructed to be similar to that here. Second, we use those utility parameters and the observed quantities X_t, q_t, and $G(\cdot)$ in the NLSY97 to recover the offer distribution, F and β through the method of simulated moments. We explicate our approach in the following sub-sections.

A. Utility Parameters

We will first pin down utility parameters using vignette survey data where respondents choose between hypothetical jobs. We capture their initial work experience and various demographic details and then show them a series of job offers meant to replicate the environment of our model. At each point they pick between the 3 hypothetical jobs and a lump sum compensation.

Here we briefly describe our vignette. For a more complete description of the instrument and sample see Section IV. In each of 14 sequential vignettes the respondents are asked to choose from a set containing three hypothetical offers and direct monetary compensation. These offers change between vignettes but every respondents sees the same set. All jobs are fully characterized by attributes in 6 dimensions including hours, variation in hours, monthly earnings, variation in pay, flexibility, and degree of advanced notice. Each dimension then has a discrete set of possible values. An example vignette is provided in Figure 1 while the options within each possible attribute are described in Table 1. As evidenced from the example vignette, the jobs are distinct contracts comparable to the $q = (W, H, z)$ of our model.

From our model in Section II, consider the problem of a worker at some intermediate point, t, in their lifetime. Returning to equation (3) and the value of a job we see that

$$V^j(\Omega_t, t) = \mathbb{E}_{W,H}\left[U(c_t, l_t, z_t) + \beta\, \mathbb{E}\left[V^j(\Omega_{t+1}, t+1)\middle|\, \Omega_t, p_t = 1\right]\right]$$

We model the vignettes with this construct in mind. We collect the inputs contained in the Ω_t, including age, demographic data, and work experience, and then ask people to choose between job options that imply different immediate rewards, $U(\cdot)$ and future discounted expected values $V^j(\Omega_{t+1}, t+1)$. These individuals repeat this process several times over changing sets of job offers. We disallow a respondent to pick their current employment status to maximize the comparability between alternatives.

The essential insight in our approach is that by fixing an individual at some time t we can use the vignette setting to trace out the utility parameters. In particular, at some time t the functional form of $V^j(\Omega_{t+1}, t+1)$ remains constant but we can still vary the inputs, Ω_{t+1}, within a controlled environment. The vignette permits us to propose a variety of possible hypothetical jobs that in turn imply a variety of future values of accumulated human capital. Unlike in typical revealed

preferences data, such as panel data, we perfectly observe all the alternatives the individual faced and can form the universe of alternative-specific values. Furthermore, we are able to control the dimensions of jobs considered such that every alternative is fully specified over the space $\mathbb{J}$ but also contains no qualities not captured in $\mathbb{J}$.

A potential concern regarding our approach is that the offered jobs represent unrealistic offers in the space of accepted, or even market-feasible offers. In Section IV we present some statistical evidence regarding observed jobs but we also note that this feature would be an advantage of stated preferences data in some situations. In particular, if we observed very little variance in hours for high-income individuals in revealed preferences-type data we might draw the conclusion that increasing income is associated with a distaste for variance in hours when it could be that varying hours for a high income individual is not cost efficient from an employer's perspective. Here, stated preference data conveniently recovers preferences over options that may be censored from observed jobs for unrelated reasons.

Another possibility is to forego stated preferences data and propose a similar set-up using revealed preference data. The primary justification for using stated preference data over revealed preference data arrives in the critical needs to observe all offered alternatives and to control for outside influences. Recovering the space of alternatives from panel data suffers from the difficulty discussed above of separating the influence of preferences from that of the offer function. An unpopular alternative may never show up in a revealed preferences data set or simply occur with low probability (potentially reflecting its unpopularity). Failure to distinguish these situations could lead to a mis-estimation of preference parameters. Furthermore, at a minimum this type of exercise requires knowing the full menu offered by any particular employment contract as well as controlling for any other relevant differences between jobs. The latter proves exceptionally challenging when working with revealed preferences as attempts to control for differences between jobs can introduce an omitted-ability bias into estimates. However, without these conditions it is difficult to construct an outside options argument, especially since actual job changes are non-random and thus potentially not indicative of preferences. The primary advantage of this vignette-type exercise, then, is that jobs are completely defined by the same six dimensions and thus there is theoretically no need to control for unobserved differences or exogenous influences.

A final relevant advantage of using vignette data is that it allows us to include individuals who

are not currently employed in a regular job (though survey respondents are paid). Consider, in particular, the agent who chooses home production every period during the work-life. Clearly, the individual may strictly prefer home production or be receiving so few offers that home production is the best alternative. By imposing sufficient structure on functional forms we could rationalize their behavior but not distinguish its source. By using stated preferences data we potentially directly identify the types of offers that would pull the individual out of home production and into the work force. Further, since our model includes the option of home production the inclusion of unemployed or out-of-the-labor-force individuals should improve estimation efficiency.

B. Offer Function Parameters

Having recovered the utility parameters using vignette survey data we use revealed preferences data to pin down the parameters underlying the offer distribution. We use the universe of observed individuals, accepted jobs, and their characteristics from the NLSY97. Knowing the preferences of workers and the resultant distribution of accepted offers allows us to make strong inferences on the primitive offer distribution. Returning to (3),

$$V^j(\Omega_t, t) = \mathbb{E}_{W,H}\left[U(c_t, l_t, z_t) + \beta\, \mathbb{E}\left[\left. V^j(\Omega_{t+1}, t+1)\right| \Omega_t, p_t = 1\right]\right],$$

and studying the properties of the offer distribution we see that the essential obstacle is to model the worker's expectations and further how they shift with accumulated work experience. This exercise requires formulating expectations over paths not observed in the data as well as for the full distribution of offers though an individual may never actually receive a large measure of those.

Our approach breaks the problem into two steps. In the first we construct an estimate of $G(\hat{q}, Z)$ using observed jobs and the set of characteristics along the six dimensions described in Table 1. This estimate informs our estimation procedure by providing the characteristics of observed jobs for different levels of work experience. This object corresponds to the offer distribution, F, that agents take as given except that it lacks unchosen alternatives. Using the estimated utility parameters we then implement a simulated method of moments procedure that recovers the parameters of the offer distribution, θ_Θ and β jointly. The procedure solves the Bellman equation in (4) by iteration using a given vector of parameters and then comparing some implied moments

of the solution to observed moments in the data. The primary moments we use for estimation center around matching the distribution of accepted offers conditional on work experience.

Section V contains the details of the data and our estimation procedure. We describe and summarize moments $\hat{G}$ as well as provide a Table detailing the moments used for estimation. We also discuss details of the estimation procedure and our chosen functional form assumptions. In the sub-section below we discuss the identification of the offer distribution given the data and process we described above.

C. Identification of the Offer Distribution

To be completed.

IV. Vignettes

The data used for estimation contains the stated preference between of 3 jobs and a lump sum monetary compensation from 14 vignettes for a set of respondents interviewed in February 2015.[6] The sample contains respondents born between 1980 and 1984 and is constructed to have an equal gender split. The survey mechanism contains machinery to prevent fraudulent responses including a fifteenth vignette that appeared identical to the others but contained hidden text instructing respondents to ignore the standing instructions and rank the jobs in a specific order.[7]

As seen in the vignette in Figure 1, a contract is a distribution of hours per week, a work schedule, a level of flexibility, and a pretax monthly pay value. The latter neatly distills variability in wage level, productivity pay, and hours into one relatable measure. In particular, the mean monthly pay is constructed as the product of mean number of hours ($\overline{H}$), the mean monthly wage rate ($\overline{W}$), and 4 weeks per month. The minimum and maximum are then equal to the limiting amount of hours the employee could face if her job varies in hours (i.e. one of the three cases in $\{(\overline{H},\overline{H}),(.8\overline{H},1.2\overline{H}),(.5\overline{H},1.5\overline{H})\}$) times the limiting wages she could receive if her pay varies (i.e. fixed or variable as $\{(\overline{W},\overline{W}),(.9\overline{W},1.1\overline{W})\}$). Referring again to Figure 1, Option A contains fixed hours and a fixed wage and thus offers no distribution in monthly pay. Option B varies in hours but not pay. Option C varies in hours and pay and contains a relatively wide variance in pay.

[6] Qualtrics, a private survey institute, commissioned the data collection over the internet on behalf of Flávio Cunha.
[7] The exact wording was "Attention filter: please choose option A to proceed.

Table 1: *Contractual Attributes*

Dimension	Possible Values	Notes	NLSY Q# (Rounds)
Advanced Notice	One week or less	You will learn your schedule each week	YEMP-WS3C (15+)
	Between 1 and 2 Weeks	You will learn your schedule 1-2 weeks before	
	3 weeks or more	You will learn your schedule at least 3 weeks in advance	
Schedule Flexibility	None.	Starting and finishing times are decided by your employer and you cannot change them on your own	YEMP-WS4B (15+)
	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	
Hours	10	hours per week	YEMP-90700 (1+)
	20	hours per week	YEMP-34428 (1+)
	30	hours per week	[Actual Value]
	40	hours per week	
Variance in Hours	Fixed	e.g. always H hours/week	
	$\pm$20%	Hours vary: Min .8$\times H$, Max 1.2$\times H$, Mean H	YEMP-WS1B (15+)
	$\pm$50%	Hours vary: Min .5$\times H$, Max 1.5$\times H$, Mean H	YEMP-WS2B (15+)
Pre-Tax Wage	Low	$0.9 \times w_{i,p} + 0.1 \times w_{i,r}$ Pessimistic	YEMP-97500 (1+)
	Mid1	$0.6 \times w_{i,p} + 0.4 \times w_{i,r}$ Low Mixture	[Actual Value]
	Mid2	$0.4 \times w_{i,r} + 0.6 \times w_{i,o}$ High Mixture	
	High	$0.1 \times w_{i,r} + 0.9 \times w_{i,o}$ Optimistic	
Variance in Pay	Fixed	Always w per hour.	YEMP-20700 (1+)
	$\pm$10%	Includes incentive pay: Min .9$\times w$, Max 1.1$\times w$, Mean w	

The space of contractual attributes has 6 dimensions each with a finite set of alternatives. A job offer is defined as some combination of these 6 contractual attributes. The values $w_{i,p}, w_{i,o}$, and $w_{i,r}$ correspond to the pessimistic, optimistic, and realistic wage offers respectively.

A notable feature of the survey's design is that it allows construction of tailored outside job offers for both employed and unemployed respondents—a novelty in the literature. Wages used in hypothetical jobs are specific to each respondent and formed by initially asking respondents to consider thirty hypothetical job offers and provide what they feel the maximum ($\overline{w}$), minimum ($\underline{w}$), and the most likely value ($\widetilde{w}$) of the wage rates would be.[8] The survey then uses four possible base wage rates constructed from the following moments:

$$\begin{aligned} w_{lo} &= 0.9 \times \underline{w}_i + 0.1 \times \widetilde{w}_i \\ w_{mid1} &= 0.6 \times \underline{w}_i + 0.4 \times \widetilde{w}_i \\ w_{mid2} &= 0.4 \times \widetilde{w}_i + 0.6 \times \overline{w}_i \\ w_{hi} &= 0.1 \times \widetilde{w}_i + 0.9 \times \overline{w}_i. \end{aligned} \tag{7}$$

[8]See Appendix A for more information on this process.

Since $\overline{w} \geq \widetilde{w} \geq \underline{w}$ is enforced during the survey it is necessarily true that $w_{lo} \leq w_{mid1} \leq w_{mid2} \leq w_{hi}$. A hypothetical job offers a distribution of monthly pretax pay if varying hours and productivity pay are included but the central tendency of said distribution will always be the product of mean hours and one of the four wage values.

The advantage of using the four ordered wage offers is considerable. First, it largely side-steps the problematic issue of non-comparability between what constitutes a high-paying or low-paying job both across people and across the pay-scale. For example, a 5% raise may or may seem like a high paying job to a respondent but a wage close to their most optimistic expected wage should necessarily seem so. Further, a typical issue with hypothetical jobs is the difficulty in constructing a reasonable value for monthly pay across individuals—as the moments of an individual's wage distribution almost certainly vary by person but likely by income group as well. The mechanism further facilitates the incorporation of the current employment status as their wage rate is also sorted into one of the four possible wage bins for estimation purposes. This technique permits the classification of their current job within the vignettes attribute structure and mitigates the effect of potentially inaccurate self-reported wages. Finally, this approach allows the incorporation of unemployed respondents for whom though their current wage is undefined but a distribution of potential offers likely exists.

Returning to Table 1, the survey looks at six dimensions within a job contract. These dimensions cover the prototypical wage and hour specification with a mean hours listing of between 10 to 40 hours per week and an hourly wage rate from on the four options constructed above. The survey also lists the amount of advanced notice the employee can expect as well as the amount of control they would have over their starting and finishing times. These last two relate to two separate issues within the variability of work hours in that even if you have perfect control over your start and finish times, the inability to forecast your work schedule imposes significant costs. Similarly, a perfectly forecastable schedule may still be less desirable than having a degree of control over starting and ending times which can potentially mitigate scheduling issues. Finally, there are two binary attributes that directly introduce unpredictability into the contracts. The first allows hours per week to follow a distribution. The second induces monthly pay to vary between plus and minus 10% each month productivity pay relative to the specified base wages.

One concern when working with stated preferences data is vulnerability to bias. The issue

Suppose you currently receive a regular monthly unemployment benefit of $800 per month. You then get three job offers from employers who would like to hire you because you have the education level, the work experience, and the qualifications that fit the requirements for a new job at their local facility. Please select which option you would most prefer going forward.

For the purposes of this exercise please assume all other aspects of the listed jobs are identical.

Option A	Option B	Option C	Current Job
HOURS: Fixed at 10 hours per week	**HOURS:** Hours vary with business volume. You will work at least 8 hours per week, at most 12 hours per week, and on average 10 hours per week.	**HOURS:** Hours vary with business volume. You will work at least 5 hours per week, at most 15 hours per week, and on average 10 hours per week.	
SCHEDULE: You will have 3 weeks or more of advance notice.	**SCHEDULE:** You will have 1 week or less of advance notice.	**SCHEDULE:** You will have 1 week or less of advance notice.	
FLEXIBILITY: A lot. You are entirely free to decide when you start and finish work.	**FLEXIBILITY:** None. Starting and finishing times are decided by your employer and you cannot change them on your own.	**FLEXIBILITY:** None. Starting and finishing times are decided by your employer and you cannot change them on your own.	**PRETAX MONTHLY PAY:** Unemployment benefit of $800 per month
PRETAX MONTHLY PAY: Fixed at $400 per month	**PRETAX MONTHLY PAY:** Varies with hours worked. Pretax monthly pay is at least $320, at most $480, and on average $400	**PRETAX MONTHLY PAY:** Varies with hours worked. Pretax monthly pay is at least $280, at most $840, and on average $560	
○	○	○	○

Figure 1: *Example vignette. Each respondent saw 12 such vignettes where they choose between 3 hypothetical jobs and a lump sum benefit. The hypothetical options were chosen using a factorial design technique so that only a subset of all possible jobs were necessary.*

arises from several possible sources. The principal suspect is an untestable hypothetical bias where individuals report being willing to pay more in a hypothetical scenario than a real one (see List (2001)) and does not have an econometric solution. This bias does not affect relative valuations between individuals nor between non-monetary attributes but can inflate the estimated valuations of non-monetary attributes in relation to earnings. Therefore, some care will be taken in the discussion of the results to account for its presence.

V. Estimation of Our Human Capital Accumulation Model

To be completed.

References

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

Blundell, Richard, Mike Brewer, and Marco Francesconi, "Job Changes and Hours Changes: Understanding the Path of Labor Supply Adjustment," *Journal of Labor Economics*, 2008, *26* (3), pp. 421–453.

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

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).

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), pp. 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), pp. 327–336.

Keane, Michael P. and Kenneth I. Wolpin, "The Career Decisions of Young Men," *Journal of Political Economy*, 1997, *105* (3), 473–522.

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 August 2014.

List, John A., "Do Explicit Warnings Eliminate the Hypothetical Bias in Elicitation Procedures? Evidence from Field Auctions for Sportscards," *American Economic Review*, December 2001, *91* (5), pp. 1498–1507.

Manning, Alan, "How Do We Know That Real Wages Are Too High?," *The Quarterly Journal of Economics*, 1995, *110* (4), 1111–1125.

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.

Rosen, Sherwin, "The Theory of Equalizing Differences," in O. Ashenfelter and R. Layard, eds., *Handbook of Labor Economics*, Vol. 1 of *Handbook of Labor Economics*, Elsevier, 1987, chapter 12, pp. pp. 641–692.

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.

A. Vignette Details

One of the goals of our survey is to measure expectations and uncertainty that individuals face about future hourly wage rates which are normally used to measure labor compensation. In the next few questions, we will ask you to report the possible range of hourly wage rates you believe you could be offered in the short run (say, in the next three months).

Q101.
Consider the following thought experiment. Imagine that in **in the next three months** you are contacted by 30 employers who would like to hire you because you have the education level, the labor market experience, and the qualifications that fit the requirements for a new job at their local facility.

Think about these 30 offers and be **pessimistic**. In your opinion, what will be the lowest hourly wage rate among these 30 offers?

Important: We are interested in the lowest hourly wage offer among these 30 employers. You may (or may not) choose to accept such offer.

. . .

Q103.
Think about these 30 offers and be **realistic**. What do you think will be the most likely value for the hourly wage rate among these 30 offers?

For your convenience, the table below reproduces your answers to the previous questions.

	Hourly wage rates
Pessimistic	$${q://QID321/ChoiceTextEntryValue}/hour
Realistic	**Q 103**
Optimistic	$${q://QID323/ChoiceTextEntryValue}/hour

Enter your answer in the space below.

Figure 2: *Respondents are prompted with the above set of instructions and then asked what the minimum, maximum, and central tendency for wage offers they expect to receive would be. Each offer is chosen seeing previous selections and controlled so as to not violate implied order relationships.*