

Interpolating Value Functions in Discrete Choice Dynamic Programming Models

Paul Sullivan*
Bureau of Labor Statistics

October 2006

Abstract

Structural discrete choice dynamic programming models have been shown to be a valuable tool for analyzing a wide range of economic behavior. A major limitation on the complexity and applicability of these models is the computational burden associated with computing the high dimensional integrals that typically characterize an agent's decision rules. This paper develops a regression based approach to interpolating value functions during the solution of dynamic programming models that alleviates this burden. This approach is suitable for use in models that incorporate unobserved state variables that are serially correlated across time and correlated across choices within a time period. The key assumption is that one unobserved state variable, or error term, in the model is distributed extreme value. Additional error terms that allow for correlation between unobservables across time or across choices within a given time period may be freely incorporated in the model. Value functions are simulated at a fraction of the state space and interpolated at the remaining points using a new regression function based on the extreme value closed form solution for the expected maxima of the value function. This regression function is well suited for use in models with large choice sets and complicated error structures. The performance of the interpolation method appears to be excellent, and it greatly reduces the computational burden of estimating the parameters of a dynamic programming model.

*The views expressed in this paper are those of the author and may not represent the views of the U.S. Bureau of Labor Statistics. Contact information: Email: Sullivan.Paul.Joseph@bls.gov. Phone: (202) 691-6593. Fax: (202) 691-6583. Bureau of Labor Statistics, Postal Square Building, Room 3105 MC 204, 2 Massachusetts Avenue N.E., Washington, D.C. 20212-0001.

1 Introduction

In recent years a large amount of research has centered on the estimation of structural discrete choice dynamic programming models of rational behavior. This type of model represents a theoretically appealing approach to modelling situations in which a forward looking agent makes decisions in the presence of uncertainty about future events. Existing research shows that dynamic programming models are a valuable tool for examining a wide range of topics spanning many different fields of economics such as industrial organization, labor, and development. Topics examined using dynamic programming models include engine replacement (Rust 1987), occupational choices (Keane and Wolpin 1997), retirement (Berkovec and Stern 1991, Rust and Phelan 1997), educational choices (Arcidiacono 2004), and decisions about marriage and cohabitation (Brien, Lillard, and Stern 2006) to name just a few examples.¹

A major impediment to the estimation of dynamic discrete choice models is the computational burden associated with solving the agent's dynamic optimization problem. Solving the optimization problem produces the value functions that are used to compute a likelihood function or set of moment conditions that are used to estimate the parameters of the model. The value functions must be computed for each feasible choice, in every time period, for every value of the state variables that influence choice-specific rewards. Computing the value functions requires evaluating the expected value of the maximum valued choice available in the next time period at each point in the state space. In general, computing the expected value of the maximum choice, called E_{max} in the remainder of this paper, involves computing a high dimensional integral. The computational burden of solving a dynamic programming model grows rapidly as the size of the state space increases, creating a well known problem known as the "curse of dimensionality" (Bellman 1957).

Several approaches have been developed to reduce the computational burden of estimating dynamic programming models. One approach takes advantage of particular functional form as-

¹See Eckstein and Wolpin (1989) for a survey of applications of discrete choice dynamic programming models.

assumptions and model structures so that evaluating high dimensional integrals is not required when solving the dynamic programming problem (Miller 1984, Rust 1987). For example, Rust (1987) assumes that the only source of randomness in his model is an extreme value error term so that Emax has a convenient closed form solution. A second approach reduces the computational burden of estimation by estimating the parameters of a dynamic programming model without solving the optimization problem (Hotz and Miller 1993, Manski 1991). Aguirregabiria and Mira (2002) build on this work and develop a method of estimating dynamic programming models that bridges the gap between the full solution and Hotz and Miller (1993) estimation approaches.² A third approach developed by Rust (1997) solves the dynamic programming problem using a method that circumvents the curse of dimensionality using randomization. Finally, Keane and Wolpin (1994) develop a method of approximately solving the dynamic programming problem that involves simulating Emax at a fraction of the state space and then interpolating Emax at the remaining points in the state space using a linear regression of Emax on the expected value functions for each choice. This approach reduces computation time because it replaces relatively slow numerical integration with comparatively fast interpolation.

This paper develops a method for solving dynamic programming problems that combines features of the interpolation method developed by Keane and Wolpin (1994) with a less restrictive version of the functional form assumptions used by Rust (1987). The key assumption is that there is one unobserved state variable, or error term, in the model that is distributed extreme value. There are no restrictions on the other error terms in the model, so additional error terms that allow for correlation across choices within a given time period or for correlation between error terms across time may be freely incorporated. Allowing unobservables to be correlated across choices and

²One limitation of these non-full solution estimation methods is that in general they cannot be used when unobservable state variables are serially correlated. See Aguirregabiria and Mira (2006) for an estimation method for dynamic games that can be used when there is permanent time invariant unobserved heterogeneity.

across time is necessary in many natural applications of dynamic programming models in order to develop economic models that are sufficiently realistic to capture important aspects of behavior. For example, any reasonable formulation of a job search model must allow for matching between workers and firms, which requires formulating and solving a dynamic programming model that incorporates serially correlated unobserved job match components. Similarly, in a dynamic model where firms make entry decisions it is desirable to allow for firm specific shocks to profitability that are correlated across markets.³

Following Keane and Wolpin (1994), the interpolation method developed in this paper involves simulating the value functions at a subset of state points and interpolating the value functions at the remaining points in the state space using a linear regression. The method departs from Keane and Wolpin (1994) by using a new interpolating regression function that takes advantage of the assumption that one of the error terms in the model is distributed extreme value.⁴ The interpolating regression involves regressing E_{\max} on a constant and the closed form solution that E_{\max} would take on if the only error term in the model was the extreme value error. The intuition behind this regression function is that the "option value" arising from the non-extreme value error terms is captured by the parameters of the interpolating regression. This regression function has the desirable theoretical property that it converges to the exact solution for E_{\max} as the standard deviations of the non-extreme value error terms approach zero. In addition, the single regressor is not vulnerable to the collinearity problems between regressors that can make it difficult to apply interpolating regressions that use individual expected value functions separately as explanatory

³In practice, only a very limited number of dynamic programming models that allow for serially correlated unobserved state variables have been estimated. See, for example, Pakes (1986), Berkovec and Stern (1991), Wolpin (1992), and Stinebrickner (2001).

⁴Since the independence of irrelevant alternatives (IIA) problem is not present in dynamic programming models, the extreme value assumption has computational benefits analogous to those of estimating a multinomial logit model as opposed to a multinomial probit model without the unappealing IIA property.

variables. This problem is likely to arise in models with large choice sets and when several choices have similar discounted expected values.

The interpolation method is tested using a dynamic structural model of occupational choices and job search. This particular application of dynamic programming techniques is well suited to examine the performance of the interpolating regression because it includes several features that are likely to be found in a wide range of applications. For example, this model incorporates a large choice set (21 choices) where the set of feasible choices varies over the state space. The model also includes a large state space and a rich error structure that allows for correlation between unobservable state variables both across choices within a given time period and across time for a given choice.

The interpolation method is evaluated by examining the extent to which the interpolated value functions generate choices that match the optimal choices produced by the model when the dynamic program is solved exactly (without using interpolation). The performance of the interpolation method is evaluated for three sets of structural parameters that create different lifecycle choice patterns and also differ in the importance of the non-extreme value error terms. The performance of the interpolation method is excellent, with at least 97% of interpolated choices matching the optimal choices across all three simulated data sets. The interpolation method is extremely accurate even when the value functions are interpolated at 99% of the state space. At this level of interpolation the value functions can be computed over 100 times faster than when the value functions are computed without using interpolation. The large decrease in computation time combined with the excellent performance of the interpolation algorithm makes estimating dynamic programming models with large choice sets, large state spaces, and complicated and realistic error structures computationally feasible.

The remainder of the paper is organized in the following manner. Section 2 discusses the solution of dynamic programming models and presents the interpolation method. Section 3 presents a model

of occupational choices and job search that is used to evaluate the performance of the interpolation method. Section 4 presents evidence on the performance of the interpolating regression, and Section 5 concludes.

2 Solving Dynamic Discrete Choice Models

This section presents a dynamic programming model and discusses the solution to the agent's optimization problem. Assume that each individual chooses an action $k_t \in D_t$ in each discrete time period, $t = 1, \dots, T$. Denote the number of elements in D_t as K_t . Alternatives are defined to be mutually exclusive. The individual's objective is to choose the feasible alternative in each time period that maximizes the discounted expected value of future rewards. The one period reward from choosing alternative k in time period t is

$$R_t(k, S_t, \psi_{kt}, e_{kt}, \varepsilon_{kt}) = U_t(k, S_t) + \psi_{kt} + e_{kt} + \varepsilon_{kt}, \quad (1)$$

where S_t is a vector of state variables that evolves deterministically over time conditional on the chosen alternative, and $U_t(k, S_t)$ is a deterministic function of the state vector. The term ψ_{kt} is a continuous random state variable that affects the reward from choosing alternative k in time period t . This variable is constant over time as long as an individual chooses action k , so $\psi_{k_{t-1}} = \psi_{kt}$ if $k_{t-1} = k_t$. Define the cumulative distribution function of ψ_{kt} as $G(\psi)$. The term e_{kt} is a continuous random state variable that may be correlated across the various alternatives within a time period but is independent across time. Define the distribution of e_{kt} as $H(e)$. The variable ε_{kt} represents a random shock to the reward from alternative k in time period t that enters the reward function additively. Assume that ε_{kt} is independent across alternatives and across time, and is distributed multivariate extreme value with variance $\pi^2\tau^2/6$. Define the cumulative distribution function of ε_{kt} as $F(\varepsilon)$. The state variables ψ , e , and ε are not observed by the econometrician, but the current

period realizations of these variables are observed by the agent when he makes his optimal choice in each period.⁵

The agent's optimization problem can be represented in terms of alternative specific value functions, $V_t(k)$. Define the expected value of the individual's optimal choice in time period $t + 1$ for an individual who is currently in time period t as

$$EW_{t+1}(S_{t+1}) = \iiint \max_{\{k \in D_{t+1}\}} \{V_{t+1}(k), k \in D_{t+1}\} dG(\psi) dH(e) dF(\varepsilon) \quad (2)$$

When computing the expected value of the optimal choice in the next time period the agent must compute the expectation over the distributions $F(\varepsilon)$, $H(e)$, and $G(\psi)$ because the time $t + 1$ realizations of these random variables are unknown in time period t . For notational simplicity the state variables ψ , e , and ε are suppressed when writing $EW_{t+1}(S_{t+1})$. In the remainder of the text the expected value $EW_{t+1}(S_{t+1})$ will be referred to as "Emax." The discounted expected value of choice k at time period t is

$$V_t(k) = R_t(k, S_t, \psi_{kt}, e_{kt}, \varepsilon_{kt}) + \delta EW_{t+1}(S_{t+1}) \text{ for } t < T, \quad (3)$$

$$V_t(k) = R_t(k, S_t, \psi_{kt}, e_{kt}, \varepsilon_{kt}) \text{ for } t = T, \quad (4)$$

where δ is the discount factor. For brevity of notation the arguments $\{S_t, \psi_{kt}\}$ are suppressed when writing the value functions. The optimization problem is solved by traditional backwards recursion techniques using equations 2-4. The backwards solution method is used to calculate $\{V_t(k), k \in D_t, t = 1, \dots, T\}$ at all possible values of S_t and ψ_{kt} .⁶

⁵The unobserved state variables ψ and e enter the reward function additively for expositional convenience, but the interpolation method may be used if these variables enter non-linearly. However, ε must enter the reward function additively.

⁶When the state space contains continuous variables they must be discretized during the solution of the model. See Stinebrickner (2000) for a detailed analysis of the issues that arise when continuous, serially correlated state variables are included in discrete choice dynamic programming models.

The major computational burden in solving the optimization problem arises from the fact that the high dimensional integral in equation 2 must be evaluated at each point in the state space. Typically this integral will not have a closed form solution, so the integral must be approximated using numerical methods such as Gaussian quadrature or Monte Carlo simulation. When the state space is large, it is extremely time consuming to numerically compute these integrals at each point in the state space during the recursive solution algorithm.

Let $\bar{V}_t(k) = V_t(k) - \varepsilon_{kt}$. A consequence of the assumption that ε_{kt} is distributed extreme value (Rust 1987) is that the expected value of the best choice next period takes the following form

$$EW_{t+1}(S_{t+1}) = \iiint \max_{\{k \in D_{t+1}\}} \{V_{t+1}(k), k \in D_{t+1}\} dG(\psi) dH(e) dF(\varepsilon) \quad (5)$$

$$= \iint \tau(\gamma + \ln[\sum_{k \in D_{t+1}} \exp(\frac{\bar{V}_{t+1}(k)}{\tau})]) dG(\psi) dH(e) \quad (6)$$

$$= \iint \Psi[\bar{V}_{t+1}(k)] dG(\psi) dH(e), \quad (7)$$

where γ is Euler's constant, and $\Psi[\bar{V}_{t+1}(k)] = \tau(\gamma + \ln[\sum_{k \in D_{t+1}} \exp(\frac{\bar{V}_{t+1}(k)}{\tau})])$. The primary benefit of the extreme value assumption is that the dimension of the integral is reduced by the number of feasible choices (K_t) because there is an ε_{kt} for each feasible choice.⁷ If the only error term in the model was ε , the expected value of the best choice would have an analytical solution and numerical integration would not be required during the solution of the model. During the solution of the model, $EW_{t+1}(S_{t+1})$ is approximated using simulation methods by averaging over R draws from the distributions $G(\psi)$ and $H(e)$,

$$\widetilde{EW}_{t+1}(S_{t+1}) = \frac{1}{R} \sum_{r=1}^R \Psi^r[\bar{V}_{t+1}(k)],$$

⁷The extreme value assumption is similarly beneficial when deriving the choice probabilities implied by the model. See Rust (1987) for the original treatment of this issue, and Berkovec and Stern (1991) for a dynamic programming model that includes an extreme value error term within a more complicated error structure. Berkovec and Stern (1991) avoid having to use simulation methods by assuming that future realizations of the non-extreme value error terms are known to the agent in the model.

where r indexes simulation draws.

2.0.1 The Keane and Wolpin Interpolation Method

Before presenting the interpolation method developed in this paper it is useful to discuss a closely related approach that is commonly used to reduce the burden of solving dynamic programming problems. Keane and Wolpin (1994) develop a method of interpolating value functions that greatly reduces the computational burden of solving and estimating dynamic programming models. Their method involves simulating $E \max$ at a small subset of state points, and then interpolating $E \max$ at the remaining points in the state space using a linear regression. The state points where $E \max$ is simulated are chosen randomly. This method speeds up the calculation of value functions because relatively slow numerical integration is replaced by comparatively fast linear interpolation at a large fraction of the state space. The general form of the regression function favored by Keane and Wolpin (1994) is

$$E \max(V_{t+1}(k)) = \max E(\bar{V}_{t+1}(k)) + \lambda \{\max E(\bar{V}_{t+1}(k)) - \bar{V}_{t+1}(k), k \in D_{t+1}\}, \quad (8)$$

where $E \max(V_{t+1}(k))$ is the expected value of the maximum, $\max E(\bar{V}_{t+1}(k))$ is the maximum of $\bar{V}_{t+1}(k)$ over the K_{t+1} feasible choices, $\max E(\bar{V}_{t+1}(k)) = \max\{\bar{V}_{t+1}(k), k = 1, \dots, K\}$, and $\lambda\{*\}$ is the functional form of the regression. Keane and Wolpin's preferred specification for the regression equation is

$$E \max - \max E = \pi_0 + \sum_{k=1}^{K_t} \pi_{1k} (\max E - \bar{V}_{t+1}(k)) + \sum_{k=1}^{K_t} \pi_{2k} (\max E - \bar{V}_{t+1}(k))^{1/2}. \quad (9)$$

Keane and Wolpin demonstrate that the interpolation error for this regression function is small using simulated data from an occupational choice model in which the agent has a maximum of four choices. The model presented in Section 3 of this paper has a much larger choice set, with a maximum of 21 possible choices. An additional complication is that the set of feasible choices

varies over the state space in this model. When the choice set varies over the state space, a separate interpolating regression must be estimated for each possible choice set. If only one regression is run, there will be a missing data problem because $\bar{V}_{t+1}(k)$ will be undefined at all points in the state space where choice k is infeasible. This is not a serious obstacle to using the Keane and Wolpin interpolation method unless the choice set varies so much over the state space that estimating multiple interpolating regressions becomes unwieldy, but it does increase the complexity of solving the optimization problem.

One potential problem that arises when using the Keane and Wolpin interpolation method that is closely related to the size of the choice set is colinearity in the regressors ($\max E - \bar{V}_{t+1}(k)$, $k = 1, \dots, K$) used in the interpolating regression. This problem frequently occurred during attempts to estimate the parameters of the model presented in Section 3 of this paper using maximum likelihood. The degree of colinearity in the regressors obviously depends on the particular set of parameter values and the fundamental structure of the model, but in general one would expect this problem to arise in any model that incorporates choices that are closely related. The model presented in section 3 of this paper includes 21 choices, and many of these choices are quite closely related because agents are allowed to combine options into dual activities. For example, the model allows workers to combine employment in each occupation with school attendance. One consequence of a high degree of colinearity in the explanatory variables of the interpolating regression is that the interpolated values of E_{\max} may be very far from the actual values of E_{\max} . This is a serious concern because any bias in the interpolated value functions translates into bias in parameter estimates. Additional problems are created by colinearity because during estimation small changes in the parameter vector may result in large changes in the interpolated value functions. The straightforward solution to the colinearity problem is to simply drop colinear regressors since the only goal of the regression is to accurately predict E_{\max} given a set of regressors. However, this solution can be quite difficult to implement within the context of estimating the parameters of a

dynamic programming problem.

The general problem is to estimate a vector of parameters, θ , by maximizing a likelihood function $L(\theta, V(\theta))$ that is a function of the parameter vector and the value functions, $V(\theta)$. The parameter vector θ , along with the structure of the model determines the degree of colinearity in the interpolating regression, but of course θ will change as the likelihood function is maximized. In practice, dropping a certain subset of colinear regressors may work very well at the initial parameter values, but there is no guarantee that dropping this same subset of regressors will work well at future iterations of the parameter vector. On a practical note, it quite difficult to write estimation code that automatically checks for colinearity and alters the explanatory variables used in the interpolating regression while the estimation program is running. In addition, changing the explanatory variables included in the interpolating regression over the course of estimation may cause convergence problems. The next section presents an interpolating regression that is not vulnerable to colinearity problems because it uses only one regressor, and also possesses desirable theoretical properties.

2.0.2 A New Interpolation Method

This paper develops an interpolation algorithm that builds on the one developed by Keane and Wolpin (1994). As in Keane and Wolpin (1994), value functions are simulated at a fraction of the state space and interpolated using a regression at the remaining points in the state space.⁸ This paper implements a new regression function that takes advantage of the assumption that the error term ε is distributed extreme value. Let k' refer to the agents optimal choice in time period t . Define $\bar{V}_{t+1}^*(k) = \bar{V}_{t+1}(k) - \psi_{kt+1} - e_{kt+1}$ for $k \neq k'$, and $\bar{V}_{t+1}^*(k) = \bar{V}_{t+1}(k) - e_{kt+1}$ for $k = k'$. Let $\Psi^*[\bar{V}_{t+1}^*(k)]$ represent the closed form solution for the Emax integral if the time $t + 1$ realizations

⁸Interpolated values of Emax are only used at points in the state space where Emax is not simulated. This implies that as the number of state points where Emax is simulated becomes large and the number of simulation draws becomes large the approximate solution approaches the exact solution.

of ψ and e that are random from the point of view of the agent are netted out of the one period utility flows,

$$\Psi^*[\bar{V}_{t+1}^*(k)] = \tau(\gamma + \ln[\sum_{k \in D_{t+1}} \exp(\frac{\bar{V}_{t+1}^*(k)}{\tau})]) \quad (10)$$

In other words, $\Psi^*[\bar{V}_{t+1}^*(k)]$ would be the exact extreme value analytical solution for Emax if ε was the only source of randomness in the model. This is not the case in this model due to the existence of the unobserved state variables ψ and e , but it suggests using the following interpolating regression based on the extreme value closed form solution for Emax,

$$\widetilde{EW}_{t+1}(S_{t+1}) = \omega_{0t} + \omega_{1t}\Psi^*[\bar{V}_{t+1}^*(k)]. \quad (11)$$

The parameters ω_{0t} and ω_{1t} are estimated by ordinary least squares, and are allowed to vary over time. The only explanatory variable in this interpolating regression is the extreme value closed form of the expected value of the maximum choice at time $t + 1$, $\Psi^*[\bar{V}_{t+1}^*(k)]$. Note that any departure of the extreme value solution ($\Psi^*[\bullet]$) for Emax and the actual simulated value of Emax ($\widetilde{EW}(\bullet)$) is due to the effects of the unobserved state variables ψ_{kt+1} and e_{kt+1} on Emax. The intuition behind this regression function is that the constant term in the regression captures the additional option value associated with ψ and e . One important advantage of this regression function is that no matter how large the choice set is, there will never be a colinearity problem in the interpolating regression because there is only one regressor. Also, this single regressor is defined at each point in the state space even when the choice set varies over the state space, so there is no need to estimate multiple interpolating regressions corresponding to each feasible choice set. In models with extremely large or complicated choice sets, this is a significant advantage over interpolation methods that use individual expected value functions separately as regressors in an interpolating regression.

Perhaps most importantly, this regression function has the desirable theoretical property that it converges to the exact solution for Emax as σ_ψ and σ_e approach 0, because $\widetilde{EW}_{t+1}(S_{t+1}) \rightarrow$

$\Psi^*[\bar{V}_{t+1}^*(k)]$ as $\sigma_\psi \rightarrow 0$ and $\sigma_e \rightarrow 0$. As these standard deviations approach zero the only source of randomness in the model is the extreme value error ε , so the regression coefficient ω_{0t} will approach zero and the coefficient ω_{1t} will approach one, and the interpolated values of E_{\max} will be exactly equal to the extreme value solution. Existing regression based approaches to interpolating value functions in dynamic programming models do not share this theoretical property. As σ_ψ and σ_e become large relative to the standard deviation of the extreme value error (σ_ε) the extreme value solution $\Psi^*[\bar{V}_{t+1}^*(k)]$ will become an increasingly poor approximation of E_{\max} because it ignores the growing option values associated with ψ_{kt} and e_{kt} , but this increasing option value will potentially be captured by the regression parameters ω_{0t} and ω_{1t} . Analysis of this interpolating regression presented in the next section indicates that this regression function performs very well across a wide range of values of σ_ψ , σ_e , and σ_ε . In addition, this regression function satisfies the theoretical restrictions on E_{\max} outlined in the Williams-Daly-Zachary Theorem (McFadden 1981).

3 A Model of Career Choices

This section presents a dynamic model of career choices that is used to evaluate the performance of the interpolating regression. The model presented in this section is a slightly simplified version of the model of occupational choices and job search developed and estimated in Sullivan (2006). The structure of the model incorporates features that are likely to be present in a wide range of applications of discrete choice dynamic programming models. The model includes a large choice set that varies over the state space, a large state space, and a relatively complicated error structure that allows the error terms in the model to be correlated across closely related choices and across time for a given choice. Although this example focuses on a particular dynamic programming model, the interpolation method developed in this paper is applicable to a general discrete choice dynamic programming model.

Each individual's career is modeled as a finite horizon, discrete time dynamic programming

problem. In each year, individuals maximize the discounted sum of expected utility by choosing between working in one of the five occupations in the economy, attending school, earning a GED, or being unemployed. Workers search for suitable wage match values across firms while employed and non-employed. Dual activities such as simultaneously working and attending school are also feasible choices. The exact set of choices available in each year depends in part on the labor force state occupied in the previous year. Each period, an individual always receives one job offer from a firm in each occupation and has the option of attending school, earning a GED, or becoming unemployed. Individuals observe all the components of the pecuniary and non-pecuniary rewards associated with each feasible choice in each decision period and then select the choice that provides the highest discounted expected utility.

3.1 Utility Function

The utility function is a choice specific function of endogenous state variables (S_t) and random utility shocks that vary over time, people, occupations, and firm matches. To index choices for the non-work alternatives, let $s = school$, $g = GED$ and $u = unemployed$. Describing working alternatives requires two indexes. Let $eq = \text{“employed in occupation } q\text{”}$, where $q = 1, \dots, 5$ indexes occupations. Also, let $nf = \text{“working at a new firm”}$, and $of = \text{“working at an old firm.”}$ Combinations of these indexes define all the feasible choices available to an individual. The description of the utility flows is simplified by defining another index that indicates whether or not a person is employed, so let $emp = \text{“employed”}$. Define the binary variable $d_t(k) = 1$ if choice combination k is chosen at time t , where k is a vector that contains a feasible combination of the choice indexes. Dual activities composed of combinations of any two activities are allowed subject to sensible restrictions such as the fact that employment and unemployment are mutually exclusive choices.

3.1.1 Choice Specific Utility Flows

This section outlines the utility flows corresponding to each possible choice. The utility flow from choice combination k is the sum of the logarithm of the wage, $w_{it}(k)$, and non-pecuniary utility, $H_{it}(k)$, that person i receives from choice combination k at time t ,

$$U_{it}(k) = w_{it}(k) + H_{it}(k). \quad (12)$$

The remainder of this section describes the structure of the wage and non-pecuniary utility flows in more detail.

2.1.1a Wages. The log-wage of worker i employed at firm j in occupation q at time t is

$$w_{it} = w_q(S_{it}) + \mu^q + \psi_{ij} + e_{ijt}. \quad (13)$$

The permanent worker-firm productivity match is represented by ψ_{ij} .⁹ This term reflects match specific factors that are unobserved by the econometrician and affect the wage of worker i at firm j . True randomness in wages is captured by e_{ijt} . This error structure allows for the job matching effects that are the foundation for the large search literature, and it also allows for wage fluctuations within jobs. In addition, the error terms in the model are correlated across choices within a given time period because each job offer consisting of a draw of ψ_{ij} and e_{ijt} may be combined with school attendance. This type of error structure can be applied in a wide range of contexts beyond matching between workers and firms. For example, in industrial organization applications one might specify a model that allows for matching between firms and particular markets (Berry 1992), or for unobserved product attributes (Berry, Levinsohn, and Pakes 1995). The detailed equations for the non-pecuniary utility flows ($H_{it}(k)$) are presented in Appendix A.

⁹It is straightforward to include person specific heterogeneity by allowing μ^q to vary across people using a discrete distribution of unobserved heterogeneity as in Heckman and Singer (1984). This approach is quite widespread in dynamic programming models, see for example Keane and Wolpin (1997).

2.1.1b Non-pecuniary Utility Flows. Non-pecuniary utility flows are composed of a deterministic function of the state vector and random utility shocks. Define $1\{\bullet\}$ as the indicator function which is equal to one if its argument is true and equal to zero otherwise. The non-pecuniary utility flow equation is

$$H_{it}(k) = [h(k, S_{it})] + \left[\phi^s 1\{s \in k\} + \phi^u 1\{u \in k\} + \sum_{q=1}^5 \phi^q 1\{eq \in k\} \right] + \varepsilon_{ikt}. \quad (14)$$

The first term in brackets represents the influence of the state vector on non-pecuniary utility flows and is discussed in more detail in Appendix A. The second term in brackets captures the effect of preferences for attending school (ϕ^s), being unemployed (ϕ^u), and being employed in occupation q (ϕ^q). The final term, ε_{ikt} , is a shock to the non-pecuniary utility that person i receives from choice combination k at time t .

3.1.2 The State Space

The endogenous state variables in the vector S_t measure human capital and the quality of the match between the worker and his current employer. Educational attainment is summarized by the number of years of high school and college completed, hs_t and col_t , and a dummy variable indicating whether or not a GED has been earned, ged_t . Possible values of completed years of high school range from 0 to 4, and the possible values of completed college range from 0 to 5, where five years of completed college represents graduate school. Let L_t be a variable that indicates a person's previous choice, where $L_t = \{1, \dots, 5\}$ refers to working in occupations one through five, $L_t = 6$ indicates attending school full time, and $L_t = 7$ indicates unemployment. Given this notation, the state vector is $S_t = \{hs_t, col_t, ged_t, L_t, \psi_t\}$.

3.2 The Optimization Problem

Individuals maximize the present discounted value of expected lifetime utility from age 16 ($t = 1$) to a known terminal age, $t = T$. At the start of his career, the individual knows the human capital wage function in each occupation, as well as the deterministic components of the utility function. Future realizations of firm specific match values (ψ 's) and time and choice specific utility shocks (ε 's and e 's) are unknown. Although future values are unknown, individuals know the distributions of these random components.

The maximization problem can be represented in terms of alternative specific value functions which give the lifetime discounted expected value of each choice for a given set of state variables, S_t . The value function for an individual with discount factor δ employed in occupation q is

$$V_t(eq, g | S_t) = U_t(eq, g | S_t) + \delta[EW_{t+1}(g | S_{t+1})^{eq}]. \quad q = 1, \dots, 5, \quad g = 0, 1 \quad (15)$$

The EW_{t+1}^{eq} terms represent the expected value of the best choice in period $t + 1$. In the remainder of the paper these expected values are referred to as "E_{max}". The expectation is taken over the random components of the choice specific utility flows, which are the random utility shocks and match values, $\{\varepsilon, e, \psi\}$.

The individual elements of the $EW_{t+1}(g | S_{t+1})^{eq}$ terms are the time $t + 1$ value functions for each feasible choice,

$$EW_{t+1}(g = 0 | S_{t+1})^{eq} = E \max\{V_{t+1}(s), V_{t+1}(u), V_{t+1}(u, g), \quad (16)$$

$$[V_{t+1}(ei, nf), V_{t+1}(m, ei, nf), m = s, g, \quad i = 1, \dots, 5],$$

$$V_{t+1}(eq, of), V_{t+1}(s, eq, of), V_{t+1}(g, eq, of) | S_{t+1}\}$$

$$EW_{t+1}(g = 1 | S_{t+1})^{eq} = E \max\{V_{t+1}(s), V_{t+1}(u), \quad (17)$$

$$[V_{t+1}(ei, nf), V_{t+1}(s, ei, nf), \quad i = 1, \dots, 5],$$

$$V_{t+1}(eq, of), V_{t+1}(s, eq, of) | S_{t+1}\}.$$

The value function for an individual who is not currently employed is

$$V_t(p \mid S_t) = U_t(p \mid S_t) + \delta EW_{t+1}(g \mid S_{t+1})^{su}, \quad p = \{s\}, \{u\}, \{u, g\}, \quad g = 0, 1 \quad (18)$$

The corresponding expected value of the maximum terms are

$$EW_{t+1}(g = 0 \mid S_{t+1})^{su} = E \max \{V_{t+1}(s), V_{t+1}(u), V_{t+1}(u, g), \\ V_{t+1}(ei, nf), V_{t+1}(m, ei, nf), \quad m = s, g, \quad i = 1, \dots, 5 \mid S_{t+1}\} \quad (19)$$

$$EW_{t+1}(g = 1 \mid S_{t+1})^{su} = E \max \{V_{t+1}(s), V_{t+1}(u), \\ V_{t+1}(ei, nf), V_{t+1}(s, ei, nf), \quad i = 1, \dots, 5 \mid S_{t+1}\}, \quad (20)$$

which consist of all feasible combinations of schooling, unemployment, and new job offers.

4 Solving the Career Decision Problem

Estimating the structural parameters of the model requires solving the optimization problem faced by agents in the model, and then using the computed value functions along with career choice data to evaluate a likelihood function or set of moment conditions. The finite horizon dynamic programming problem is solved using standard backwards recursion techniques.

4.1 Solution and Interpolation

This section discusses the solution method for the dynamic programming problem, and discusses how interpolation can be used to solve the optimization problem more quickly.

4.1.1 Distributional Assumptions

Assume that firm specific match values and randomness in wages are distributed i.i.d normal: $\psi_{ij} \sim N(0, \sigma_\psi^2)$, and $e_{ijt} \sim N(0, \sigma_e^2)$. The firm specific match values are part of the state space because the value function associated with a job depends on the wage match value (ψ_{ij}) for worker

i at firm j . The distribution of this variable is continuous, which causes a problem because the state space becomes infinitely large when a continuous variable is included. This problem is solved by using a discrete approximation to the distribution of wage match values (ψ_{ij}) when solving the value functions. See Stinebrickner (2000) for a detailed analysis of different solutions to the problems that arise when serially correlated error terms are included in a dynamic discrete choice model.

Assume that the random choice-specific utility shocks are distributed extreme value, with distribution function $F(\varepsilon) = \exp\{-\exp(-\frac{\varepsilon}{\tau})\}$, and with variance $\pi^2\tau^2/6$. The assumption that the ε 's are distributed extreme value simplifies the computation of the value functions and choice probabilities. See Rust (1987,1997) for examples of dynamic programming models which assume that the only error term in the model is distributed extreme value. Note that adopting the extreme value assumption does not result in the independence of irrelevant alternatives (IIA) problem in a dynamic setting even if the extreme value error is the only source of randomness in the model.¹⁰ In a dynamic model, the extreme value assumption provides substantial computational benefits without the drawbacks associated with estimating a static multinomial logit model as opposed to a static multinomial probit model.

4.1.2 Calculating the Value Functions

Computing the value functions requires calculating the value of every feasible choice at each point in the state space over the agent's entire time horizon. This requires evaluating the high dimensional integrals found in equations 16 through 20. To illustrate what is involved in computing E_{\max} ,

¹⁰See Rust (2005) for a discussion of this issue.

consider computing the following expected value,

$$EW_{t+1}(g = 0 \mid S_{t+1})^{eq} = E \max\{V_{t+1}(s), V_{t+1}(u), V_{t+1}(u, g), \quad (21)$$

$$[V_{t+1}(ei, nf), V_{t+1}(m, ei, nf), m = s, g, i = 1, \dots, 5],$$

$$V_{t+1}(eq, of), V_{t+1}(s, eq, of), V_{t+1}(g, eq, of) \mid S_{t+1}\}.$$

$$= E \max\{V_{t+1}(k), k \in D_{t+1}(g = 0)^{eq} \mid S_{t+1}\}, \quad (22)$$

where $D_{t+1}(g = 0)^{eq}$ is the set of feasible choices at time $t + 1$. The expectation is taken over the random error terms corresponding to each of the 21 feasible choices at time $t + 1$. The utility flow equations shown in Section 3 reveal that there are a total of 32 time $t + 1$ realizations of the unobserved state variables whose values are unknown to the agent in time period t . More specifically, there are 21 non-pecuniary utility shocks $(\varepsilon_{i1}, \dots, \varepsilon_{i21})$ corresponding to each choice, 5 match values $(\psi_{i1}, \dots, \psi_{i5})$ and 5 wage shocks (e_{i1}, \dots, e_{i5}) for new job offers, and one wage shock for the agent's current job (e_{i6}) .¹¹ The realizations of these error terms at time $t + 1$ are unknown to the agent at time t , but the agent knows the distributions of these error terms. The expected value is the following 32 dimensional integral

$$EW_{t+1}(g = 0 \mid S_{t+1})^{eq} = \int \int \int \max\{V_{t+1}(k), k \in D_{t+1}(g = 0)^{eq} \mid S_{t+1}\} dF(\varepsilon_{i1}, \dots, \varepsilon_{i21}) \quad (23)$$

$$dG(\psi_{i1}, \dots, \psi_{i5}) dH(e_{i1}, \dots, e_{i6}).$$

As discussed in Section 2, the assumption that the random utility shock, ε , is distributed extreme value implies that this expected value takes the following form

$$EW_{t+1}(g = 0 \mid S_{t+1})^{eq} = \int \int \Psi[\bar{V}_{t+1}(k)] dG(\psi_{i1}, \dots, \psi_{i5}) dH(e_{i1}, \dots, e_{i6}). \quad (24)$$

The integral shown in Equation 24 can be simulated in a straightforward manner using random

¹¹The time $t + 1$ realization of ψ for the agent's current job is not random from the point of view of the agent because ψ_{ij} is constant over the duration of the match between worker i and firm j .

draws from the distributions $G(\psi)$ and $H(e)$,

$$\widetilde{EW}_{t+1}(g = 0 \mid S_{t+1})^{eq} = \frac{1}{R} \sum_{r=1}^R \Psi^r[\bar{V}_{t+1}(k)], \quad (25)$$

where r indexes simulation draws. The other Emax terms can also be simulated in this manner. Throughout this paper 100 draws from the joint distribution of the errors are used to simulate Emax. Antithetic acceleration is used to reduce the variance of the simulated integrals. See Geweke (1988) for a discussion of antithetic acceleration, and Stern (1997) for a survey of the applications of simulation methods in the economics literature.

Simulating the integral found in equation 25 at each point at the state space can be so time consuming that estimation of the parameters of the model becomes impractical. The interpolating regression developed in this paper suggests evaluating the simulated integrals at a fraction of the state space in each time period and interpolating the remaining points using the regression function presented in section 2 of this paper,

$$\begin{aligned} \widetilde{EW}_{t+1}(g = 0 \mid S_{t+1})^{eq} &= \omega_{0t} + \omega_{1t}\tau(\gamma + \ln[\sum_{k \in D_t} \exp(\frac{\bar{V}_{t+1}^*(k)}{\tau})]) \\ &= \omega_{0t} + \omega_{1t}\Psi^*[\bar{V}_{t+1}^*(k)]. \end{aligned} \quad (26)$$

All $\widetilde{EW}_{t+1}(\ast)$ terms shown in equations 14-19 are interpolated using this type of regression function.

5 Assessing the Performance of the Interpolation Method

This section evaluates the performance of the interpolation method presented in the previous section by comparing the solution to the model obtained using interpolation to the exact solution to the model.

5.1 Simulated Data

One way of assessing the performance of the interpolation method is to compare simulated career choice data generated by the model when the value functions are computed at each point in the

state space to simulated data from the model when the value functions are interpolated. Any differences between the simulated choices generated from the full solution and interpolated value functions reflect approximation error caused by interpolation.¹²

It is straightforward to use the structural model to simulate choices. First, given a vector of parameters the value functions are computed (either exactly or using interpolation) at each point in the state space. At this point, the value functions for the agent are known up to a draw of the error terms in the model. Choices are simulated by drawing a complete set of errors for the agent in time period one that reflect the values of random wage shocks (e 's), job match values (ψ 's), and utility shocks (ε 's) for each feasible choice. The simulated choice is simply the one with the highest value. The simulated choice is used to update the state vector, and choices are then simulated in this manner, moving forward in time, for the agents entire time horizon. The model is used to generate simulated choice sequences for 2,000 simulated people over a 10 year time horizon.

The interpolated and exact simulated choices are computed for three versions of the parameter vector which generate significantly different choice sequences. A complete description of the parameter values used to generate each of the three simulated data sets is found in Appendix B. A key difference between the three versions of the parameter vector is the importance of the various error terms in determining career choices. The parameter values used to generate data set 1 imply a large role for the extreme value non-pecuniary utility shock relative to job matching and randomness in wages. The standard deviation of the job match value (σ_ψ) is .275, the standard deviation of the random wage shock (σ_e) is .306, and the standard deviation of the random utility shock (σ_ε) is

¹²The terms "full solution" and "exact solution" are used interchangeably to refer to solving the dynamic programming problem by simulating the value functions at all points in the state space, and not using any interpolation. More accurately, the simulation solution method is still an approximate solution for a fixed number of simulation draws because the simulated Emax's converge to the true Emax's as the number of simulation draws approaches infinity. However, experimentation shows that increasing the number of simulation draws beyond 100 leads to essentially no change in the simulated Emax's.

4.14. These parameters are reasonable values given the structural estimates of a closely related occupational choice model found in Sullivan (2006). Given these parameter values, the extreme value closed form solution for E_{\max} should be a relatively close approximation to the true value of E_{\max} , so it seems likely that the interpolating regression based on the extreme value solution for E_{\max} will perform well.

In the second simulated data set the standard deviations are changed so that the extreme value error is a less important determinant of career choices by using the following error standard deviations: $\sigma_{\psi} = .275$, $\sigma_e = 3.45$, and $\sigma_{\varepsilon} = 4.14$. Note that the total amount of randomness in the model increases from data set 1 ($\sigma_{\psi} + \sigma_e + \sigma_{\varepsilon} = 4.72$) to data set 2 ($\sigma_{\psi} + \sigma_e + \sigma_{\varepsilon} = 7.86$), and the importance of the random wage shock relative to the random non-pecuniary utility shock is much larger in data set 2. The large increase in σ_e from .306 to 3.45 increases the departure of the actual solution for E_{\max} from the extreme value closed form solution that is used as a regressor in the interpolating regression. Comparing the interpolated and exact choices in the second data set shows how well the regression function based on the extreme value solution for E_{\max} performs when $\Psi^*[\bar{V}_{t+1}^*(k)]$ is a poor approximation to E_{\max} . The third data set is generated using the following error standard deviations: $\sigma_{\psi} = .275$, $\sigma_e = 2$, and $\sigma_{\varepsilon} = 4.14$. In addition, several utility flow parameters are changed so that the simulated choices in the third data set are quite different from the choices found in data sets 1 and 2.

The occupational choice patterns in the three simulated data sets are summarized in Figure 1. The simulated choices in data sets 1 and 2 exhibit qualitatively similar choice-age profiles. One noteworthy feature of the first two data sets is the large upward age trend in professional employment, with 65% of the workers in data set 1 working as professionals at age 25, and 55% of the workers in data set 2 working as professionals at age 25. Across all three simulated data sets the proportion of people attending school declines sharply with age. This happens because the investment value of schooling declines with age and because the consumption value of schooling

declines with age. The simulated choices in data set 3 are quite different from those in the first two data sets. Some of the largest differences are that unemployment is quite common in the third data set and professional employment is very rare compared to the first two data sets. In addition, employment as craftsmen and laborers is much more common in the third data set compared to the first two data sets. Overall there is a more even choice distribution in the third data set compared to the first two data sets.

5.2 Results: Assessing the Performance of the Interpolation Method

Tables 1-4 present evidence on the performance of the interpolation method in each of the three simulated data sets. Table 1 shows the proportion of simulated choices in the data generated using interpolated value functions that match the choices found in the data generated without using interpolation. Any differences between these simulated choices are the result of approximation error in the interpolated value functions. For example, the first entry in the first column of Table 1 indicates that when E_{max} is simulated at 5% of the state space and interpolated at the remaining 95% of the state space, 99.9% of the choices in time period 1 in the simulated data match the optimal choices when interpolation is not used. The percentage of correct choices is quite stable across time periods, with an overall match rate of 99.6% across all time periods when 90% of the state space is interpolated. The final entry in the first column of Table 1 shows that the computer program that interpolates E_{max} at 95% of the state space runs 43 times faster than the program that simulates E_{max} at all points in the state space. The amount of time it takes to solve the dynamic programming problem is reduced from 6 hours to only 8 minutes when the value functions are interpolated at 95% of the state space. Given that the dynamic programming problem must be solved repeatedly during the estimation of the parameters of a dynamic discrete choice model as a likelihood function is maximized or a set of moment conditions are evaluated, using interpolation

greatly expands the scope of dynamic programming problems that are feasible to estimate.¹³

One striking feature of Table 1 is that the performance of the interpolating regression does not appear to deteriorate as the number of state points used in the interpolating regression decreases. The percentage of correct choices is constant at 99.6% when the value functions are interpolated at 95%, 99%, or 99.8% of the state space. Of course, the amount of time it takes to solve the dynamic programming problem decreases as the level of interpolation increases because an increasing number of relatively slow integral simulations are replaced by fast regression interpolations. Interpolating the value functions at 99% of the state space reduces computation time by a factor of 114, while interpolating the value functions at 99.8% of the state space reduces computation time by a factor of 819.

The results shown in Table 1 indicate that the interpolating regression provides extremely accurate predictions of E_{\max} even when the value functions are interpolated at a very large fraction of the state space. In some respects this is not surprising because in data set 1 the standard deviations σ_ψ and σ_e are relatively small, so one could do fairly well by interpolating E_{\max} with the extreme value solution $\Psi^*[\bar{V}_{t+1}^*(k)]$ without even estimating an interpolating regression. In other words, the coefficient estimates for the interpolating regression are quite close to $\omega_{0t} = 0$ and $\omega_{1t} = 1$ in many time periods for this set of parameter values. The error standard deviations used to generate data set 1 are based on maximum likelihood estimates from the closely related occupational choice model estimated in Sullivan (2006), so they should be considered reasonable parameter values.¹⁴ These parameter estimates imply that there is more dispersion in random

¹³In most applications of dynamic programming models the computational burden of evaluating a likelihood function or set of moment conditions during estimation is insignificant relative to the computational burden of solving the dynamic programming problem.

¹⁴The major differences between the models presented in Sullivan (2006) and the one presented in this paper are: 1) For simplicity, this version of the model rules out within firm occupational mobility, 2) The model in Sullivan (2006) allows for person-specific unobserved heterogeneity in occupation specific ability (μ 's) and preferences (ϕ 's) using the Heckman Singer (1984) approach. Neither change in the model is likely to impact the performance of the

shocks to non-pecuniary utility than in random wage shocks or job match values.

Table 2 examines whether the performance of the interpolating regression is robust to changes in the standard deviations of the error terms in the model. Specifically, data set 2 is generated using parameter values that increase the importance of random wage shocks in determining career choices relative to the first data set. Most importantly, the error standard deviations are increased in such a way that the extreme value solution for $E_{\max}(\Psi^*[\bar{V}_{t+1}^*(k)])$ used as the interpolating regressor is a poorer approximation of the actual value of E_{\max} . As σ_ψ and σ_e increase, $(\Psi^*[\bar{V}_{t+1}^*(k)])$ becomes a poor approximation of E_{\max} because it does not capture the increasing option value of mobility associated with large values of σ_ψ and σ_e . Data set 2 reveals the extent to which the regression parameters ω_{0t} and ω_{1t} are able to capture the departure of E_{\max} from $\Psi^*[\bar{V}_{t+1}^*(k)]$ when the importance of the non-extreme value error terms is large relative to the importance of the extreme value error.

The results shown in Table 2 demonstrate that the interpolating regression performs extremely well when the standard deviation of the normally distributed random wage shock is large relative to the extreme value utility shock. The interpolating regression captures the additional option value associated with an increase in the variance of the wage shocks through the constant in the interpolating regression. The R-squared for the interpolating regression remained at approximately .99 in each time period. When the value functions are interpolated at 95% of the state space 97.3% of the choices generated using interpolated value functions match the true choices. The match rate declines slightly over time from 98.5% in time period 1 to 95.1% in time period 10. This happens because errors in simulated choices in one time period create errors in the state variables that affect the value of future choices. For example, suppose that interpolation error causes a choice in year 1 for a simulated person to be attending school instead of unemployment. Attending school affects wages and utility differently across occupations, so in later years the simulated person is more likely

interpolating algorithm.

to choose to work in occupations where education is highly rewarded.

As in data set 1, the performance of the interpolating does not appear to deteriorate as the value functions are interpolated at an increasingly large fraction of the state space. Interpolating the value functions at 99.98% of the state space allows the dynamic programming problem to be solved 819 times faster than when the value functions are simulated across the entire state space, but the interpolated value functions still lead to the optimal choice 97.6% of the time. The overall match rate of approximately 97% in data set 2 is only slightly lower than the 99% match rate found in data set 1. Overall, the performance of the interpolating regression using the second set of parameter values is quite encouraging. The interpolating regression performs extremely well even when the extreme value solution for E_{max} is far from the true value of E_{max} .

Further information about the performance of the interpolating regression in data set 2 is shown in Table 3, which presents the average percent absolute deviation of the interpolated value functions from the actual value functions for the second data set. The entries in the table are computed using the following formula

$$\%Abs. Dev. = 100 \times \frac{|V(interpolated) - V(exact)|}{|V(exact)|},$$

averaged over all the value functions found in the simulated data. On average, the interpolated and exact value functions differ by only .32% when the value functions are interpolated at 99.98% of the state space, so the interpolation method is extremely accurate. The magnitude of the percent absolute deviation is virtually constant across all levels of interpolation, so these statistics are not reported here for the 95% and 99% interpolation levels. The excellent performance of the interpolation method according to this metric is not surprising because the extremely close match between the choices generated using the interpolated and exact value functions shown in Table 2 is only possible if the interpolation method produces value functions that are extremely close to their "true" values.

These results show that using the extreme value closed form solution for $E_{\max}(\Psi^*[\bar{V}_{t+1}^*(k)])$ in an interpolating regression performs extremely well even when the error standard deviations are set at values such that $\Psi^*[\bar{V}_{t+1}^*(k)]$ is a poor approximation of E_{\max} . The intuition behind this result is that the interpolating regression captures the departure of E_{\max} from $\Psi^*[\bar{V}_{t+1}^*(k)]$ through the constant of the interpolating regression. As the option value of job search increases moving from data set 1 to data set 2, the constant in the interpolating regression also increases.

Table 4 presents evidence on the performance of the interpolation method in data set 3. The third data set demonstrates substantially different lifecycle choice patterns compared to the first two data sets, as shown in Figure 1. The major difference is that workers are more evenly distributed across choices than in the first two data sets. This implies that the value functions for the various choices are closer to each other in magnitude in the third data set compared to the first two data sets, so there is more room for small interpolation errors to result in differences between the choices generated using the interpolated and exact value functions. The results shown in Table 4 show that the interpolation method performs quite well in the third data set, with approximately 98% of the interpolated choices matching the exact choices. As in the previous two data sets the performance of the interpolation method is very stable as the fraction of the value functions that are interpolated increases.

6 Conclusion

This paper presents a method of approximately solving dynamic discrete choice models that builds on the simulation and interpolation method developed by Keane and Wolpin (1994). The key assumption is that one of the error terms (unobserved state variables) in the model is distributed extreme value. Other error terms may be freely incorporated to allow for various types of correlation between the error terms associated with different choices in a given time period or between choices across time. Value functions are simulated at a subset of state points and interpolated at the

remaining points in the state space using a new regression function based on the extreme value closed form solution for the expected maxima of the value function.

One advantage of this approach is that when the standard deviations of the non-extreme value error terms are small the interpolated value functions are guaranteed to be close to their actual values since the interpolating regression function converges to the exact solution for E_{\max} as the standard deviations of the non-extreme value errors approach zero. On a more practical note, the regression is based on a single regressor so it avoids the colinearity problems that can arise when using interpolation methods that use individual expected value functions as regressors. Experience shows that these problems are likely to arise across a wide range of parameter values in models with large choice sets, especially when the structure of the model tends to produce choices that have similar discounted expected values. In addition, the single regressor is defined at each point in the state space even if the choice set varies over the state space, so one does not need to estimate multiple interpolating regressions corresponding to each feasible choice set.

The performance of the interpolating regression is evaluated using a dynamic model of career choices that incorporates a large choice set and an error structure that incorporates correlation across time due to job matching and correlation across choices within a given time period. The performance of the interpolation method is evaluated using three sets of parameter values that differ in the relative importance of the extreme value and non-extreme value errors and generate different optimal choice sequences. Across all simulated data sets the interpolated value functions generate choice sequences that are extremely close to the optimal choice sequences generated by exact solution to the value functions. Evidence is also presented that the differences between the interpolated and actual value functions are extremely small on average, on the order of a fraction of a percentage point.

Overall, the evidence presented in this paper suggests that the interpolating regression based on the extreme value solution for E_{\max} performs extremely well even when the extreme value solution

is a poor approximation to the actual value of E_{\max} . This method appears to be well suited for use in models with large choice sets and complicated error structures. The large decrease in the amount of computation time associated with using interpolation to solve the dynamic programming problem makes it possible to estimate increasingly realistic, and complicated, dynamic programming models.

Figure 1: Simulated Career Choices in Three Data Sets

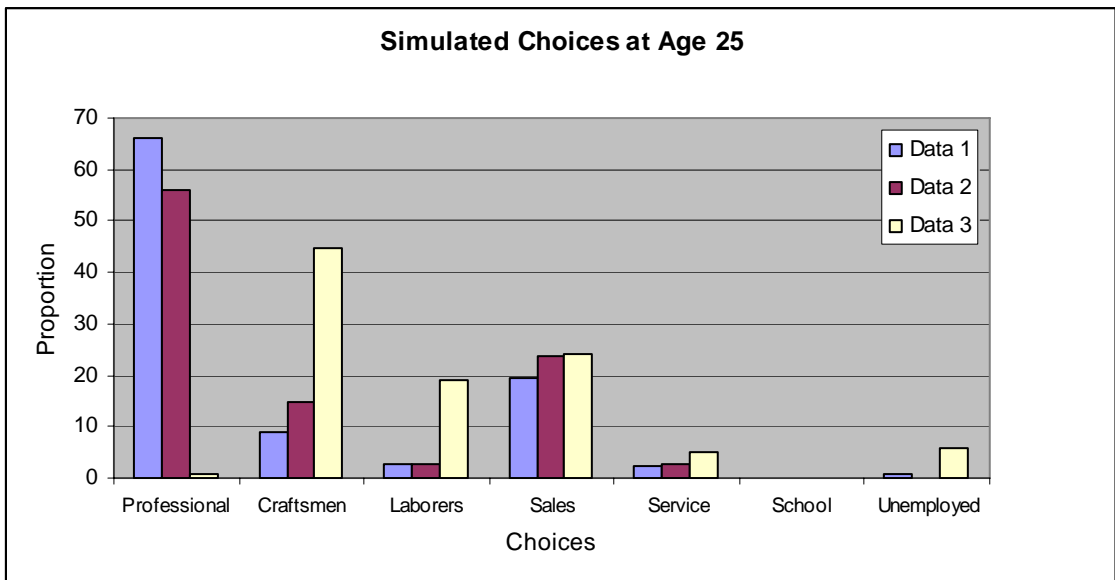




Table 1
Proportion Correct Choices – Data Set 1

	% of State Space Interpolated		
	95%	99%	99.98%
Time Period			
1	.999	.999	1.000
4	.998	.999	.998
7	.994	.995	.993
10	.993	.993	.992
All	.996	.996	.996
Number of times faster than exact solution	43	114	819

Notes: Based on a simulated sample of 2,000 people. Emax is simulated using 100 draws of the errors and antithetic acceleration is used to reduce the variance of the simulated integrals. Entries represent the proportion of simulated choices generated using different levels of interpolation that match the simulated choices when the model is solved exactly (without interpolation).

Table 2
Proportion Correct Choices – Data Set 2

	% of State Space Interpolated		
	95%	99%	99.98%
Time Period			
1	.985	.984	.983
4	.976	.978	.980
7	.964	.968	.968
10	.951	.958	.958
All	.973	.976	.976
Number of times faster than exact solution	43	114	819

Notes: See notes for Table 1.

Table 3
Percent Difference Between Simulated and Actual Value Functions – Data Set 2

% of State Space Interpolated	99.98%
Time Period	
1	.37%
4	.14%
7	.35%
10	.65%
All	.32%

Notes: Based on a simulated sample of 2,000 people. Emax is simulated using 100 draws of the errors and antithetic acceleration is used to reduce the variance of the simulated integrals. Entries represent the percent absolute deviation of the simulated value functions from the actual value functions.

Table 4
Proportion Correct Choices – Data Set 3

	% of State Space Interpolated		
	95%	99%	99.98%
Time Period			
1	.992	.997	.995
4	.982	.981	.981
7	.986	.983	.984
10	.988	.986	.986
All	.986	.986	.985
Number of times faster than exact solution	43	114	819

Notes: See notes for Table 1.

References

- [1] Aguirregabiria, Victor and Pedro Mira (2002). "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models." *Econometrica*, v. 70, no. 4.
- [2] Aguirregabiria, Victor and Pedro Mira (2006). "Sequential Estimation of Dynamic Discrete Games." forthcoming in *Econometrica*.
- [3] Arcidiacono, Peter (2004). "Ability Sorting and the Returns to College Major." *Journal of Econometrics*, v. 121: 343-375.
- [4] Bellman, Richard (1957). "Dynamic Programming." Princeton University Press.
- [5] Berkovec, James, and Steven Stern (1991). "Job Exit Behavior of Older Men." *Econometrica*, v. 59, no. 1: 189-210.
- [6] Berry, Steven (1992). "Estimation of a Model of Entry in the Airline Industry." *Econometrica*, v. 60, no. 4: 889-917.
- [7] Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium." *Econometrica*, v. 63, no. 4: 841-890.
- [8] Brien, Michael, Lee Lillard and Steven Stern (2006). "Cohabitation, Marriage and Divorce in a Model of Match Quality." *International Economic Review*, v. 47, no. 2.
- [9] Eckstein, Zvi and Kenneth Wolpin (1989). "The Specification and Estimation of Discrete Choice Dynamic Programming Models." *Journal of Human Resources*, v. 24: 562-598.
- [10] Geweke, John (1988). "Antithetic Acceleration of Monte Carlo Integration in Bayesian Inference." *Journal of Econometrics*, v. 38: 73-89.

- [11] Heckman, James, and Burton Singer (1984). "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data." *Econometrica*, v. 52: 271-320.
- [12] Hotz, Joseph and Robert Miller (1993). "Conditional Choice Probabilities and the Estimation of Dynamic Programming Models." *Review of Economics and Statistics*, v. 60: 497-530.
- [13] Keane, Michael and Kenneth Wolpin (1994). "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence." *The Review of Economics and Statistics*, v. 76, no. 4: 648-672.
- [14] Keane, Michael, and Kenneth Wolpin (1997). "The Career Decisions of Young Men." *Journal of Political Economy*, v. 105 : 474-521.
- [15] Manski, Charles. "Nonparametric Estimation of Expectations in the Analysis of Discrete Choice Under Uncertainty." In *Nonparametric and Semiparametric Methods in Econometrics and Statistics*, edited by W. Barnett, J. Powell, and G. Tauchen. pp. 259-276. MIT Press, Cambridge.
- [16] McFadden, Daniel (1981). "Econometric Models of Probabilistic Choice." In C. Manski and D. McFadden (eds.) *Structural Analysis of Discrete Data with Econometric Applications*, pp. 198-272. MIT Press, Cambridge.
- [17] Miller, Robert A. (1984). "Job Matching and Occupational Choice." *Journal of Political Economy*, v. 92: 1086-1120.
- [18] Pakes, Ariel (1986). "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks." *Econometrica*, v. 54, no. 4.
- [19] Rust, John (1987). "Optimal Replacement of GMC Bus Engines: An Empirical Model of Howard Zurcher." *Econometrica*, v. 55: 999-1033.

- [20] Rust, John (1997). "Using Randomization to Break the Curse of Dimensionality." *Econometrica*, v. 65, no. 3.
- [21] Rust, John (2005). "Structural Estimation of Markov Decision Processes." In *Handbook of Econometrics Vol. 4*, edited by R. Engle and D. McFadden. North-Holland.
- [22] Rust, John and Christopher Phelan (1997). "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets." *Econometrica*, v. 65: 781-831.
- [23] Stern, Steven (1997). "Simulation Based Estimation." *Journal of Economic Literature*, v. 35, no. 4.
- [24] Stinebrickner, Todd (2000). "Serially Correlated Variables in Dynamic Discrete Choice Models." *Journal of Applied Econometrics*, v. 15: 595-624.
- [25] Stinebrickner, Todd (2001). "A Dynamic Model of Teacher Labor Supply." *Journal of Labor Economics*, v. 19, no. 1.
- [26] Sullivan, Paul (2006). "A Dynamic Analysis of Educational Attainment, Occupational Choices, and Job Search." Working paper.
- [27] Wolpin, Kenneth (1992). "The Determinants of Black-White Differences in Early Employment Careers: Search, Layoffs, Quits, and Endogenous Wage Growth." *Journal of Political Economy*, v. 100: 535-560.

Appendix A: Utility Flow Equations

The remaining portion of the non-pecuniary utility function contains the non-pecuniary employment and non-employment utility flows along with the schooling cost function. This utility flow equation is specified as

$$\begin{aligned}
 h(k, S_{it}) &= \left[\sum_{q=1}^5 \theta_q(S_{it}) 1\{eq \in k\} \right] \\
 &+ C^s(S_{it}) 1\{s \in k, emp \notin k\} + C^{sw}(S_{it}) 1\{s \in k, emp \in k\} \\
 &+ b(S_{it}) 1\{u \in k\} + C^g(S_{it}) 1\{g \in k\}.
 \end{aligned} \tag{27}$$

The term in brackets contains the occupation and firm specific non-pecuniary utility flows. The occupation specific portion of this flow, $\theta_q(S_{it})$, is a function of the state vector that is allowed to vary over occupations. The second line of equation 27 contains the schooling cost functions for attending school while not employed ($C^s(S_{it})$) and employed ($C^{sw}(S_{it})$). The final components of the non-pecuniary utility flow are the deterministic portions of the value of leisure enjoyed while unemployed, $b(S_{it})$, and the cost function for earning a GED, $C^g(S_{it})$.

The deterministic portion of the occupation specific human capital wage function is

$$\begin{aligned}
 w_q(S_{it}) &= \beta_1^q age_{it} + \beta_2^q age_{it}^2/100 + \beta_3^q hs_{it} + \beta_4^q col_{it} + \beta_5^q 1[age_{it} \leq 17] + \\
 &\beta_6^q 1[age_{it} \geq 18 \cap age_{it} \leq 21] + \beta_7^q ged_{it}
 \end{aligned} \tag{28}$$

Let NF_t be a dummy variable indicating whether or not the individual is in his first year of employment at a firm after being employed at a different firm in the previous period. Let hd_t and cd_t represent dummy variables that indicate receipt of a high school or college diploma. The non-pecuniary utility flow equation for occupation q is

$$\begin{aligned}
 \theta_q(S_{it}) &= \alpha_1^q age_{it} + \alpha_2^q age_{it}^2/100 + \alpha_3^q (hs_{it} + col_{it}) + \alpha_6^q hd_{it} \\
 &+ \alpha_7^q cd_{it} + \alpha_8^q ged_{it} + \alpha_9^q 1[Lit > 5] + \alpha_{10}^q NF_{it} \quad q = 1, \dots, 5.
 \end{aligned} \tag{29}$$

The cost function for attending school is

$$\begin{aligned}
c^S(S_{it}) &= \gamma_{s1}age_{it} + \gamma_{s2}age_{it}^2/100 + \gamma_{s3}hd_{it} + \gamma_{s4}cd_{it} + \gamma_{s5}hs_{it} + \gamma_{s6}col_{it} + \gamma_{s7}1[L_{it} \neq 6] \\
c^{SW}(S_{it}) &= \gamma_{sw1}age_{it} + \gamma_{sw2}age_{it}^2/100 + \gamma_{sw3}hs_{it} + \gamma_{sw4}col_{it} + \gamma_{s7}1[L_{it} \neq 6] \\
&\quad + \gamma_{sw6}(hs_{it} \leq 4) + \gamma_{sw7}(hs_{it} = 4 \cap col_{it} \leq 4) + \gamma_{sw8}(col_{it} \geq 4).
\end{aligned} \tag{30}$$

The deterministic portion of the unemployment utility flow, $b(S_{it})$, is set equal to zero because the non-wage utility flow coefficients are only identified relative to a base choice, as in any discrete choice model.

The final utility flow equation represents the utility derived from earning a GED. The deterministic portion of the GED utility flow is

$$c^g(S_{it}) = \gamma_{g1} + \gamma_{g2}age_{it}. \tag{31}$$

Appendix B
Structural Parameter Values – Wage Equation: Dataset 1

<i>Variable</i>	<u>Occupations</u>				
	<i>Professional & managers</i>	<i>Craftsmen</i>	<i>Operatives & laborers</i>	<i>Sales & clerical</i>	<i>Service</i>
<u>Log Wage Equation:</u>					
Age (β_1)	-.018	.096	.003	.036	-.011
Age ² /100 (β_2)	.089	-.408	.036	-.036	.205
Years of high school (β_3)	.044	.013	.054	.029	.020
Years of college (β_4)	.097	.046	.031	.073	.097
Age \leq 17 (β_5)	-.272	-.069	-.201	-.180	-.032
18 \leq Age \leq 21 (β_6)	-.272	-.036	-.165	-.193	-.042
GED (β_7)	.020	.001	.055	.021	.011
<u>Error Standard Deviations</u>					
True randomness in wages (σ_ϵ)	.306				
Pecuniary firm match value (σ_ψ)	.275				
Extreme value parameter (σ_ϵ)	4.14				

Structural Parameter Values - Intercepts: Dataset 1

<i>Variable</i>	<u>Occupations</u>	
	<u>Log-wage Intercepts</u>	<u>Non-pecuniary Intercepts</u>
	$(\mu's)$	$(\phi's)$
Professional & managerial	9.677	-28.8
Craftsmen	9.108	-21.0
Operatives & laborers	9.346	-14.3
Sales & clerical	9.319	-22.8
Service	9.162	-19.2
School	--	6.8

Structural Parameter Values – Non-Pecuniary Utility: Dataset 1

<i>Variable</i>	<i>Parameter</i>
Discount factor (δ)	.95
<u>School Utility Flow</u>	
Age (γ_{s1})	-3.666
Age ² /100 (γ_{s2})	9.591
Attending college (γ_{s3})	.671
Attending graduate school (γ_{s4})	-2.264
Years of high school (γ_{s5})	.569
Years of college (γ_{s6})	.488
<u>School While Employed Utility Flow</u>	
Age (γ_{sw1})	-5.271
Age ² /100 (γ_{sw2})	24.74
Years of high school (γ_{sw3})	4.138
Years of college (γ_{sw4})	1.054
<u>GED Utility Flow</u>	
Constant (γ_{g1})	-.950
Age (γ_{g2})	-10.409
<u>Switching Costs</u>	
Cost of moving to a new firm (firm to firm transitions) (α_{10})	2.661
School re-entry cost (γ_{s7})	-2.376
Cost of moving to a new job from non-employment (α_9)	2.658
<u>Costs of Working while Attending School</u>	
Work in high school (γ_{sw6})	6.497
Work in college (γ_{sw7})	11.548
Work in graduate school (γ_{sw8})	12.093

Structural Parameter Values – Non-Pecuniary Utility: Dataset 1

Occupations

<i>Variable</i>	<i>Professional & Managers</i>	<i>Craftsmen</i>	<i>Operatives & Laborers</i>	<i>Sales & Clerical</i>	<i>Service</i>
<u>Employment Non-Pecuniary Utility Flows:</u>					
Age (α_1)	1.927	2.035	.860	1.761	.850
Age ² /100 (α_2)	-7.995	-10.098	-4.105	-10.689	-4.028
Education (α_3)	.773	-.649	-.620	.248	.024
High school diploma (α_6)	.639	2.222	1.749	1.862	.756
College diploma (α_7)	2.492	4.803	4.319	5.127	3.527
GED (α_8)	1.422	1.718	2.335	1.711	2.982

Dataset 2

Changes to the Parameter Vector, Relative to Dataset 1

Error Standard Deviations

True randomness in wages (σ_e)	3.45 (.306 in Dataset 1)
---	-----------------------------

Note: All other parameters unchanged from Dataset 1.

Dataset 3
Changes to the Parameter Vector, Relative to Dataset 1

<i>Variable</i>	
<u><i>Log-wage Intercepts</i></u>	<u><i>Non-pecuniary Intercepts</i> (ϕ's)</u>
Professional & managerial	-31.8 (-28.8)
Craftsmen	-24.0 (-21.0)
Operatives & laborers	-12.3 (-14.3)
Sales & clerical	-26.6 (-22.8)
Service	-16.2 (-19.2)
School	-3.8 (6.8)
Non-Pecuniary age effect for professionals (α_1)	.800 (1.927)
True randomness in wages (σ_e)	2.00 (.306)

Note: Parameter value in dataset 1 in parentheses. All other parameters unchanged from Dataset 1.