

The Specification and Estimation of Dynamic Stochastic Discrete Choice Models

A Survey

Zvi Eckstein
Kenneth I. Wolpin

I. Introduction

This paper is a survey of a rapidly growing literature on methods for solving and estimating dynamic stochastic discrete choice models. The major characteristics of this literature are: (1) the optimization problem is forward looking and contains some stochastic elements; (2) the choice set only takes on discrete values; (3) the approach is structural, that is the parameters to be estimated come from objective functions (tastes and/or technology) and constraints; and (4) the econometrics is closely connected to the theory, and specifically, the error structure is motivated in the problem as an integral part of the optimization.

The importance and usefulness of discrete choice models is now well established and such models are routinely applied. While it has always seemed that dynamic optimization under uncertainty would provide a better description (and prediction) of behavior, until recently methods did not exist to consistently (structurally) incorporate dynamic elements into discrete choice models. It is this task that the new literature addresses. The major tool for these new methods is dynamic programming (Bellman 1957). Dynamic programming is a recursive solution method for optimization problems which have a dynamic structure, and can be ap-

plied with a discrete, continuous or mixed discrete-continuous choice set. The major insight of dynamic programming is that the solution of a multidimensional problem can be reduced to a recursive solution of a sequence of two period problems.

To illustrate the idea of dynamic programming and, more importantly, to illustrate the contribution of the literature we survey, consider the very simple problem of when to consume an indivisible good given an endowment of the good. Let $d(t) = 1$ if the good is consumed at time t , $d(t) = 0$ otherwise, and $R(d(1), \dots, d(T))$ be the reward or utility from consuming the good at the T possible times. The direct solution of the problem would require that we make T comparisons of utility, i.e., $R(1, 0, \dots, 0)$, $R(0, 1, \dots, 0)$, \dots , $R(0, 0, \dots, 1, 0)$, $R(0, 0, \dots, 0, 1)$. However, suppose we proceed as follows. If by the beginning of period T the good has not been consumed the only possibility is to consume it at T , so let $V(T) = R(0, 0, \dots, 0, 1)$ which is the maximal lifetime utility at T given the state at T . Suppose now that we have reached $T - 1$ without consuming the good. If we consume it at $T - 1$ then utility is $R(0, 0, \dots, 1, 0) = R(T - 1)$, while if we do not then utility is $R(0, 0, \dots, 0, 1) = V(T)$. Define maximal lifetime utility at $T - 1$ as $V(T - 1) = \max(R(T - 1), V(T))$, which is the maximum of utility if we consume the good at $T - 1$ and maximal utility if we do not. Similarly if we have reached $T - 2$ without consuming the good, then if we don't consume it at $T - 2$ we obtain $V(T - 1)$, while if we do, we obtain $R(0, 0, \dots, 1, 0, 0) = R(T - 2)$. Thus our maximal utility at $T - 2$ is $V(T - 2) = \max(R(T - 2), V(T - 1))$. In general, then we can continue to work backwards recursively making pairwise comparisons, with maximal utility at any $t < T$ obeying the recursive equations $V(t) = \max(R(t), V(t + 1))$. The dimensionality of the problem has been reduced from a T -wise comparison to $T - 1$ pairwise comparisons.

This example illustrates a very simple optimal stopping problem. The dynamics arise because the decision at time t explicitly depends on prior decisions, namely we can only consume the good now if we did not previously consume it. The individual decides whether or not to consume the good at t , given that it has not been consumed up to t , according to whether $R(t) \geq V(t + 1)$. One might think of $V(t + 1)$ as the reservation reward or utility at $t + 1$, i.e., the reward if one postpones consumption of the good.

The above model predicts that if all individuals have the same reward function, they will all make the same decision. Suppose we have data on when a sample of observably homogeneous individuals each consumed the good. It is quite likely that the data will reveal variation in the timing of consumption, in which case the model would clearly be rejected. One way to allow for heterogeneity in observed behavior is to assume that the

Zvi Eckstein is professor of economics at Tel-Aviv University and Boston University.

Kenneth I. Wolpin is a professor of economics at the University of Minnesota. The authors would like to thank Fusun Gonul, Hide Ichimura, Chuck Manski, Bob Miller, Ariel Pakes, John Rust, Steven Stern, and the editors for their useful comments on an earlier draft of this survey.

reward function is stochastic, namely given by $R(t) + \epsilon(t)$, and to assume that individuals draw at each t from the density function for $\epsilon(t)$.¹ The analogue of the lifetime utility function under uncertainty is $V(t) = \max\{R(t) + \epsilon(t), E[V(t+1)|\Omega(t)]\}$ where $\Omega(t)$ is the information set at t and $E[V(t+1)|\Omega(t)]$ is the conditional expectation of $V(t+1)$ taken over all future ϵ 's. Now, under the further simplifying assumption that the $\epsilon(t)$'s are i.i.d. over time, the model predicts that the proportion of individuals who choose to consume the good at t (given the good has not yet been consumed prior to t) is given by the probability that each of their draws of $\epsilon(t)$ exceeded $E[V(t+1)|\Omega(t)] - R(t)$. The assumption that $\epsilon(t)$ is known to the individual but not to us makes the observed decision random from our perspective. A comparison of these predicted probabilities over time to the corresponding observed individual choices or sample proportions forms the basis for the estimation of the reward function. It is this extension of the use of dynamic programming to estimation that is the contribution of this new literature.

We begin the paper in Section II with a general specification of a dynamic stochastic discrete choice model that encompasses all existing studies in this literature. We show how the general framework nests, in order of presentation, a job search model (Wolpin 1987), a patent renewal model (Pakes 1986), an engine replacement model (Rust 1987), an armed services retention model (Gottz and McCall 1987), a labor force participation model (Eckstein and Wolpin 1986, Gottz 1987), a retirement model (Berkovc and Stern 1987, Rust 1987), a fertility model (Wolpin 1984, Montgomery 1988, Hotz and Miller 1989), and a job matching model (Miller 1984).² Section III discusses alternative solution methods for these types of dynamic programming problems. We explicitly demonstrate solution methods for several of the preceding models. In Section IV we explain the maximum likelihood method of estimation as it is applied to these models and briefly touch on an alternative estimation method, the method of simulated moments. Section V briefly surveys two semi-reduced form approaches to estimating the structural parameters recently advanced in Hotz and Miller (1988) and in Manski (1988 and 1990 forthcoming). We also suggest several alternative approximation methods. A standard search model is used as an example in each of the sections, so that the reader can follow a familiar single model from its initial specification through its estimation. The final section summarizes the paper and

concludes with a few additional topics in labor economics that can be explored using the methods discussed here.

II. A General Model

We consider a general model of I discrete choices over T discrete periods of time, where T is either finite or infinite. In each period an individual chooses one of the I possible alternatives (e.g., employment; status, occupation, etc.), where the indicator $d_i(t) = 1$ if alternative i is chosen at time t and $d_i(t) = 0$ otherwise, that is, if alternative i is not chosen by the agent. Alternatives are mutually exclusive, i.e., $\sum d_i(t) = 1$.³ The objective of the individual at any time t , $t = 0, 1, \dots, T$, is to maximize

$$(1) \quad E \left[\sum_{j=1}^T \beta^{j-t} \sum_{i \in I} R_i(j) d_i(j) | \Omega(t) \right]$$

where $0 < \beta < 1$ is the individual's discount rate, $E(\cdot)$ is the mathematical expectations operator, $\Omega(t)$ is the individual's information set at time t which includes all past and current realizations of the variables that directly or indirectly affect the value of Equation (1) and $R_i(t)$ is a random variable representing the individual's reward if alternative i is chosen at time t . $R_i(s)$ for $s \leq t$ obviously belongs to the individual's information set at time t , $\Omega(t)$.

Maximization of (1) is accomplished by choice of the optimal sequence of control variables $\{d_i(t)\}_{i \in I}$ for $t = 0, 1, \dots, T$ which are functions of information that is available when the decision is made. Define the maximal expected value of the reward at time t

$$(2) \quad V(\Omega(t)) = \sup_{\{d_i(t)\}_{i \in I}} E \left[\sum_{j=t}^T \beta^{j-t} R(j) | \Omega(t) \right]$$

where $R(j) = \sum_{i \in I} R_i(j) d_i(j)$ is the actual reward at time j . The function V depends only on the information set at time t , and obeys the dynamic programming equation⁴

$$(3) \quad V(\Omega(t)) = \max_{i \in I} \{L_i V(\Omega(t))\}$$

where L_i is the alternative-specific operator defined by

$$(4) \quad L_i V(\Omega(t)) = R_i(t) + \beta E[V(\Omega(t+1)) | d_i(t) = 1], \quad t = 0, \dots, T.$$

1. The additivity of the disturbance in the current return function is not necessary.
 2. We begin with the search model because it is the example we use to illustrate the solution and estimation methods throughout the paper. The rest of the ordering is for pedagogical reasons, in part due to the complexity of the models, and does not indicate the relative importance of the papers.
 3. Alternatives can always be redefined to satisfy this assumption.
 4. The sup operator is equal to the max operator if the maximum exists.

The dynamics of the problem are due to the dependence of the function V at time $t + 1$ on the choice of d_i , $i = 1, \dots, I$, at time t and possibly before.

This general discrete choice problem describes all of the different structurally estimated dynamic discrete choice models found in the literature.⁵ We now turn to a description of these models.

A. Optimal Stopping Models

Optimal stopping models are special cases of discrete choice dynamic programming models. We shall now show how the models of optimal stopping that exist in the literature fit the above general model. The most familiar optimal stopping model in economics is that of job search (Lippman and McCall 1976, Mortensen 1970).

1. A Job Search Model (Wolpin 1987)

Wolpin (1987) structurally estimated a standard two-state job search model in which a wealth-maximizing individual faces a known wage offer distribution, a constant cost of search, and a known per-period probability of receiving a wage offer. In each period of a finite horizon the individual decides whether to accept an offer if one is received. If an offer is rejected or if none is received the individual continues to search. Rejected offers cannot be recalled.

To place the search model in the framework of the previous section, note first that there are two alternatives, $I = 2$. Let $d_1(t) = 1$ if the individual is not employed (searching) and $d_1(t) = 0$ otherwise, and $d_2(t) = 1$ if the individual is employed and $d_2(t) = 0$ otherwise.⁶ The reward function is given by

$$(5a) \quad R_1(t) = b,$$

$$(5b) \quad R_2(t) = \begin{cases} w(t) & \text{if } d_1(t-1) = 1 \\ R_2(t-1) & \text{if } d_2(t-1) = 1 \end{cases}$$

where b is net income if the individual searches (unemployment compensation minus search costs) and $w(t)$ is the wage at time t . If the individual is not employed at t , a wage offer is drawn with a known probability p

5. A formal treatment of this type of problem exists, for example, in Whittle (1985), as well as in other texts on dynamic stochastic programming.

6. In the case of $I = 2$ (two alternatives) it is more convenient to use one indicator because $d_1(t) + d_2(t) = 1$. But, for the general case ($I > 2$) the general notation is less cumbersome and so we employ it throughout the paper.

from a time-independent distribution function $F(w(t))$ which is also known.⁷ With probability $1 - p$ no wage offer is received and net income is equal to b .

The search model described above is an example of an optimal stopping problem, i.e., once an offer is accepted the individual will never return to the nonemployment state. The stopping property results from the assumption in Equation (5b) that no new wage offers are received after a job is accepted as indicated by the fact that the reward at t is the previous period wage if the individual is employed at $t = 1$.⁸ Properties of the solution both in the finite and infinite horizon cases are presented in Lippman and McCall (1976) and the economics literature contains numerous extensions of the basic model.

2. Patent Renewal (Pakes 1986)

Pakes formulated and estimated an optimal stopping model of patent renewal. Each period over a finite horizon a patentee must decide whether or not to pay an annual renewal fee in order to keep the patent in force. If the renewal fee is not paid then the patent is permanently canceled.

In terms of the general model there are two choice variables ($I = 2$), $d_1(t) = 1$ if the patent is not renewed and $d_1(t) = 0$ otherwise, and $d_2(t) = 1$ if the patent is renewed and $d_2(t) = 0$ otherwise. The reward for each choice is given by

$$(6a) \quad R_1(t) = 0,$$

$$(6b) \quad R_2(t) = \begin{cases} \max\{\delta R_2(t-1), z(t)\} - c(t) & \text{with probability } 1 - e^{-\sigma R_2(t-1)} \\ 0 & \text{with probability } e^{-\sigma R_2(t-1)} \end{cases}$$

$c(t)$ is the renewal cost of the patent at time t , $\theta > 0$ and $0 < \delta < 1$ are known parameters, $z(t)$ is a random variable with density function $g(z, t) = \sigma(t)^{-1} \exp\{-(\gamma + z/\sigma(t))\}$ and $\sigma(t) = \phi^{t-1}\sigma$, where σ , γ and ϕ are known parameters. In addition $R_2(1)$ is assumed lognormal with mean μ and standard deviation σ_R . Pakes allows this initial draw to be different for each firm. The idea is that with some probability which depends on the previous period return, $e^{-\sigma R_2(t-1)}$, the patent is determined to have no value, while with probability $1 - e^{-\sigma R_2(t-1)}$ the patent either depreciates in value or a new use is found and the patent value is $z(t)$.

7. Wolpin actually assumes that p is duration dependent in a known way, but for expositional purposes we ignore this minor complication.

8. For an example of a model in which offers are received while on the job, see Burdett (1978). See Flinn and Heckman (1982) for a discussion of search models in a continuous time setting.

The patent model, like the search model, is an optimal stopping problem because of the institutional restriction that a lapsed patent cannot be renewed, i.e., $R_2(t+1) = 0$ if $d_1(t) = 1$.⁹ Pakes's model is different from the job search model of the previous section mainly due to the assumed serial correlation in the reward function, $R_2(t)$. That is, the reward from the renewed patent at time $t-1$, $R_2(t-1)$, affects the distribution from which the return on the patent is determined at time t , while the distribution of the wage, $w(t)$, in the search model, was assumed independent of past realizations and decisions. This specification by Pakes has important implications for the solution and estimation of the model.

3. Engine Replacement (Rust 1987)

Rust models the optimal replacement of bus engines. As in the previous model there are two alternatives ($I = 2$), $d_1(t) = 1$ if the decision is not to replace an engine, $d_1(t) = 0$ otherwise, and $d_2(t) = 1$ if the decision is to replace the engine, $d_2(t) = 0$ otherwise. The reward function is given by

$$(7a) \quad R_1(t) = -c(x(t)) + \epsilon_1(t),$$

$$(7b) \quad R_2(t) = -RC - c(0) + \epsilon_2(t)$$

where $c(x(t))$ is the operating cost of a bus having an engine with $x(t)$ miles, RC is the replacement cost and $\epsilon_i(t)$ ($i = 1, 2$) are random variables that affect (additively) the reward function $R_i(t)$ and are known to the agent but are not observed by the econometrician. Rust assumed that Harold Zurcher, the superintendent of maintenance at the Madison, Wisconsin Metropolitan Bus Company, decides about engine replacement by maximizing (1) with $T = \infty$ using the reward function (7) and the law of motion for $x(t)$ specified below.

To complete the model, Rust formulated the evolution of engine mileage as a nondecreasing stochastic process governed by a Markovian transition probability law

$$(8) \quad P(x(t+1)|x(t), d_1(t), d_2(t), \theta_2) = \begin{cases} \theta_2 \exp(\theta_2(x(t+1) - x(t))) & \text{if } d_1(t) = 1 \text{ and } x(t+1) \geq x(t) \\ \theta_2 \exp(\theta_2 x(t+1)) & \text{if } d_2(t) = 1 \text{ and } x(t+1) \geq 0 \\ 0 & \text{otherwise} \end{cases}$$

where θ_2 is a parameter.

The engine replacement problem is also an optimal stopping problem.

9. However, unlike the search model, because of the stochastic nature of the reward function, in particular because of z , an agent might wish to renew a lapsed patent.

i.e., once an engine is replaced it will never be re-used as long as the same structure applies, because mileage, $x(t)$, is nondecreasing over time and the replacement cost $c(x(t))$, is increasing in $x(t)$.

Rust generalized the model to I decision variables. Two critical assumptions characterize his specification for the general case as well as for the case of $I = 2$:

(i) $\epsilon_i(t)$, $i = 1, \dots, I$, is additive in the reward function of each alternative i and each $\epsilon_i(t)$ has an identical and independent over time extreme value distribution.

(ii) The transition density function of the vector of observed variables $(x(t))$ is conditionally independent of $\epsilon_i(t)$, $i = 1, \dots, I$. (Rust defines this as a Conditional Independence Assumption, since it does not require that $\epsilon_i(t+1)$ be independent of $x(t+1)$ and $x(t+1)$ may depend on $x(t)$); $p(x(t+1), \epsilon(t+1)|x(t), \epsilon(t)) = q(\epsilon(t+1)|x(t+1))p(x(t+1)|x(t))$, where $\epsilon(t+1)$ is a vector over the I alternatives.¹⁰

Unlike Pakes's model where the unobservable is serially correlated, here only the observed variables can be serially correlated. This assumption is critical for the algorithm that Rust developed to solve the model.

4. Air Force Retentions (Gottz and McCall 1986)

Gottz and McCall formulate and estimate a dynamic model of the stay/leave decision confronting officers in the U.S. Air Force. The essential features of the model, using our prior notation with $d_1(t) = 1$ if the officer leaves and zero otherwise and $d_2(t) = 1$ if the officer stays and zero otherwise, are:

$$(9a) \quad R_1(t) = r(t)(m_k(t) - a_k(t)) + w(t)$$

$$(9b) \quad R_2(t) = m_k(t) + z + \epsilon_k(t)$$

where $m_k(t)$ is the level of military pay at grade level k with t years of service inclusive of basic allowances $a_k(t)$, $r(t)$ is the fraction of net (of allowances) military pay received upon retiring from the service ($r(t) = 0$ for $t < 20$), $w(t)$ is the civilian wage received upon retirement, z is an individual-specific nonpecuniary reward to military service and $\epsilon_k(t)$ is z transient nonpecuniary reward. In addition an officer faces a promotion transition probability matrix, $P_{kj}(t)$ which reflects the probability of being promoted from grade level k to j after t periods of service. There is a terminal period given by the mandatory separation years of service which is grade-level specific. Note that the terminal period is itself random because future grade levels are not known with certainty. Gottz and McCall assume that z is distributed as extreme-value and ϵ is i.i.d. normal.

10. Independence implies conditional independence.

B. Sequential Choice Models

In the class of sequential choice models the individual's optimal decision may be characterized by switching between activities all of which, in contrast to optimal stopping models, may have been previously chosen.

1. A Labor Force Participation (LFP) Model (Eckstein and Wolpin 1986 and Gonul, forthcoming).

Suppose that the individual can choose between two alternative employment states ($l = 2$): $d_1(t) = 1$ if the individual does not work in the labor market, $d_1(t) = 0$ otherwise, and $d_2(t) = 1$ if the individual works in the labor market, $d_2(t) = 0$ otherwise.¹¹

The reward function is:

$$(10a) \quad R_1(t) = \alpha_1 + Y(t),$$

$$(10b) \quad R_2(t) = \alpha_2 + w(t) + Y(t)$$

where $\alpha_1 > 0$ is the monetary equivalent of the psychic value of home time, $Y(t)$ is nonearned income, $\alpha_2 < 0$ is the monetary equivalent of the psychic disutility of work, and $w(t)$ is the wage received while working. The wage is a random variable assumed to depend positively on accumulated market work experience $K(t)$. Specifically, Eckstein and Wolpin assumed that the wage was log-linear in experience,

$$(11) \quad \ln w(t) = \beta_0 + \beta_1 K(t - 1) + \epsilon(t)$$

where the law of motion for experience is given by

$$(12) \quad K(t) = K(t - 1) + d_2(t).$$

Eckstein and Wolpin further assumed that $\epsilon(t)$ is a normally distributed random variable with zero mean and a constant variance (σ^2). They included other exogenous variables, e.g., schooling, in reward function (10) and in the wage Equation (11).

The main dynamic aspect of the model arises because of the effect of experience ($K(t)$) on the wage, and therefore on future work decisions. The difference between the job search model and the labor force participation model is in the structure of wage offers. In the latter a new independent wage (conditional on experience) is drawn regardless of whether an offer has already been accepted. In the job search model no new wage offer is forthcoming if the individual had previously accepted an offer. Hence, the labor force participation model allows for workers to alternate

between work states depending upon the current wage offer and taking into account the effect of their current decision on future wages.

2. A Retirement Model (Berkovec and Stern 1988; and Rust 1988)

Berkovec and Stern estimate a retirement model which allows for three alternatives ($l = 3$): $d_1(t) = 1$ if the individual works full time at an old job, zero otherwise, $d_2(t) = 1$ if the individual works part time at an old job, zero otherwise, $d_3(t) = 1$ if the individual works part time at a new job, zero otherwise, and $d_4(t) = 1$ if the individual retires. Thus, all alternatives are feasible at each period; for example, an individual cannot choose to work full (part) time at an old job if the individual was not already working full (part) time. The reward functions are:

$$(13a) \quad R_i(t) = \bar{W}_i(t, K_i, X) - cJ + e_i(t) + \mu_i \quad \text{for } i = 1, 2,$$

$$(13b) \quad R_i(t) = \bar{W}_i(t, K_i, X) - cJ + e_i(t) + \mu_i \quad \text{for } i = 3, 4,$$

$$(13c) \quad R_5(t) = \bar{W}_R(t, X) + e_R(t) + \mu_R$$

where \bar{W}_i and \bar{W}_R are the deterministic component of wages when employed full time and part time, respectively, K_1 and K_2 are the number of prior periods worked full time or part time, respectively, on the current job (by definition $K_3 = K_4 = 0$), $X(t)$ are other exogenous variables such as schooling, health, race, and age, J is an indicator function equal to one if $K_i = 0$, for all i , and zero otherwise, reflecting job mobility costs, \bar{W}_R is the monetary equivalent of the psychic value of leisure, e_i , e_R , and e_X are i.i.d. random components of the rewards, and μ_i , μ_R , and μ_X are individual specific permanent components of the reward. These last components introduce unobserved correlation in the reward over time, i.e., serial correlation in the aggregate disturbance. The random errors (the e_i 's) are assumed to be distributed as extreme value while the μ_i 's are assumed to be normal. Berkovec and Stern assume the state variables, in particular health, to be perfectly foreseen.¹² Except for some exogenously imposed permanent retirement age, transitions into and out of retirement are feasible, i.e., "temporary" retirement is potentially an optimal state.

3. A Fertility Model (Wolpin 1984; Montgomery 1988; and Hoiz and Miller 1989)

In Wolpin's (1984) fertility model the individual faces two alternatives, $d_1(t) = 1$ if the decision is to have a child, $d_1(t) = 0$ otherwise, and $d_2(t) =$

11. We adopt the framework in Eckstein and Wolpin rather than that in Gonul because it is simpler to expose.

12. Note that Rust's retirement model allows the state variables to follow a Markovian transition probability matrix.

If the decision is not to have a child, and $d_2(t) = 0$ otherwise. Letting the stock of children at time t be (for simplicity abstracting here from the mortality of children assumed by Wolpin)

$$(14) \quad M(t) = \sum_{s=0}^t d_1(s),$$

Wolpin's specification can be written as

$$(15a) \quad R_1(t) = U(M(t-1) + 1, x(t)) + \epsilon(t)(M(t-1) + 1),$$

$$(15b) \quad R_2(t) = U(M(t-1), x(t)) + \epsilon(t)M(t-1),$$

where $x(t)$ is the level of consumption at time t . The dynamics arise because today's fertility decision affects the current stock of children which, because children are durables, affects future utility. The main difference between this model and all of the models previously described is that the error, $\epsilon(t)$, is the same for both choices and is proportional to the state variable, in this case the stock of children. Wolpin assumed that $\epsilon(t)$ is normal and independent over time.¹³

Holtz and Miller model the joint labor supply and fertility decision. The fertility component includes three alternatives, $d_1(t) = 1$ if the couple contracepts (imperfectly), $d_2(t) = 1$ if the couple does not contracept, and $d_3(t) = 1$ if the couple sterilizes. The conception probability depends explicitly on contraceptive decisions; sterility is irreversible. The reward functions are:

$$(16a) \quad R_1(t) = U(x(t), l(t), M(t-1)) + d_1(t)\epsilon_{1t},$$

$$(16b) \quad R_2(t) = U(x(t), l(t), M(t-1)) + d_2(t)\epsilon_{2t},$$

$$(16c) \quad R_3(t) = U(x(t), l(t), M(t-1)) + d_3(t)\epsilon_3(t).$$

$x(t)$ is family consumption and l_t is the mother's leisure time. $M(t-1)$ is the service flow from the existing stock of children which, unlike in Wolpin, is assumed to depend on the age distribution of the existing children. Specifically, $M(t-1) = \sum_{i=1}^{t-1} a_i b_{t-i} + a_{12} \sum_{i>12} b_{t-i}$, where $b_{t-i} = 1$ if a child is born during period $t-i$. Holtz and Miller further assume that existing offspring require maternal time, where the amount of the time devoted to a child is age-specific. Goods inputs to children are assumed to depend only on the stock of children, independent of their ages. The

13. Wolpin also allowed for a permanent individual-specific unobservable which could take on two possible values, i.e., individuals could be of two types in terms of their marginal utility of children. See Heckman and Singer (1984) for a discussion of methods of introducing individual-specific unobserved heterogeneity.

woman faces a stochastic wage which is only imperfectly forecastable. The ϵ_{it} 's ($i = 1, 2, 3$) are assumed to be independently and identically distributed over time and contraceptive choices. The joint distribution at any time t is assumed to be multivariate extreme value.¹⁴ It should be noted that the additional complexity evident in this model requires a different estimation method than the other models presented in this section. Discussion of their method is deferred to Section V.

4. Job Matching (Miller 1984) and Brand Choice (Eckstein, Horisky, and Raban 1988)

Miller (1984) extended and estimated the job matching model of Jovanovic (1979). Eckstein, Horisky, and Raban (1988) used the same structure for estimating a dynamic brand choice model. Both assumed an infinite horizon so that the optimal decision rule is stationary. The decision variable is $d_i(t)$ with $d_i(t) = 1$ if the individual chooses alternative i and $d_i(t) = 0$ otherwise, and $i \in I$. The reward to the individual for choosing alternative i (a job within an occupation or a brand of a product) out of I alternatives is

$$(17) \quad R_i(t) = \psi(t) + \xi_i + \sigma_i \epsilon_i(t).$$

$\psi(t)$ is a deterministic reward that may depend on time, is known with certainty and is common to all $i \in I$; ξ_i is a match-specific random reward that distinguishes choice i (job or brand) from the other alternatives, but is not directly observed by the individual either before or after i is chosen. Note that the reward function for a given alternative is unaffected by the choice of any other alternative. The individual has a prior on the distribution of ξ_i assumed to be normal with mean γ_i and standard deviation σ_i . The last term is an unobserved random shock to the reward where $\sigma_i > 0$ is known and $\epsilon_i(t)$ has a standard normal distribution and is independent and identical for all t and i . The individual observes $R_i(t)$ and $\psi(t)$ and tries to learn about ξ_i , the true reward associated with choice i . The dynamics arise because of this learning, i.e., remaining with a firm provides not only an immediate reward but also information about the future reward in the firm.

Since $\psi(t)$ is known and common to all i , the decision depends on the observed difference $R_i(t) - \psi(t)$. Given the independence among alternatives and over time, beliefs about the reward from the i th choice can be characterized as a conditional normal distribution, $N(\gamma_i(t), \delta_i(t))$, where

14. Montgomery also estimates a contraceptive choice model. His formulation is a direct application of Rust's framework.

the Bayesian (e.g., DeGroot 1970) updating rule implies that

$$(18) \quad \gamma_t(t) = \left[\delta_t^{-2} \gamma_t + \sigma_t^{-2} \sum_{s=0}^{t-1} (R_t(s) - \psi(s)) d_t(s) \right]$$

$$\times (\delta_t^{-2} + M_t(t-1) \sigma_t^{-2})^{-1},$$

$$\delta_t(t) = (\delta_t^{-2} + M_t(t-1) \sigma_t^{-2})^{-1/2}$$

where

$$M_t(t-1) = \sum_{s=0}^{t-1} d_t(s).$$

In Miller's model two jobs are assumed to be in the same occupation if and only if they have the same values for γ_t , δ_t , and σ_t , and there are assumed to be a large number of jobs in any occupation. In each period the individual chooses an occupation from which to sample a job, drawing a match-specific value ξ_t from the distribution $N(\gamma_t, \sigma_t)$ and a random shock $\sigma_t \epsilon_t(t)$. The longer the individual remains on that job, the more information is obtained about ξ_t . Note that in (18) the variance is declining in $M_t(t-1)$, thus inducing a nonstationarity. Only one job can be sampled in any period and there is no recall. In the special case where $\sigma_t = 0$ and $\psi(t) = 0$, $R_t(t) = \xi_t$, and because $R_t(t)$ is observed, learning is completed in one period. In this case, the optimal policy follows an optimal stopping rule as in the job search context; that is, once a job is maintained past one period (after learning is complete) it will be maintained forever.

C. A Brief Summary

From a methodological point of view, the important difference between these models is in their stochastic structure. As we shall see, the choice of error structure is crucial for the way these models are solved and estimated. In the table below we present a summary of the models described above in terms of the independence (I) or nonindependence (N) of the error term over several dimensions. The only special feature of the table is that we separate individual-specific unobserved heterogeneity (Column 3) from serial correlation (Column 2). The reader should note that this table does not indicate necessarily that the methods employed in the particular models are limited to the error structure imposed, although that may be true in some cases. The actual limitations will be addressed in later sections.

Table 1
Summary of Stochastic Assumptions

	Time or Age	Alternatives	Individuals ^a
Wolpin (search)	I	I	I
Pakes	N	I	I
Goiz and McCall	I	I	N
Rust	I	I	I
Eckstein and Wolpin	I	I	N
Gonul	I	I	I
Berkovec and Stern	I	N	N
Wolpin (fertility)	I	N	N
Hotz and Miller	I	I	N
Montgomery	I	I	N
Miller	N	I	I
Eckstein, Horsky, and Raban	N	I	I

a. Nonindependence takes the form of unobserved permanent individual-specific heterogeneity.

III. Solution Methods

A general characterization of the optimal plan is possible for some models without the existence of a closed form solution. However, in order to estimate the parameters of any of the models presented in the previous section we have to obtain a closed form (not necessarily analytical) characterization of the solution for the optimization problem (1). This task has occupied applied mathematicians for many years and is still an active research area. The actual computation of the solution may be technically different for finite horizon and infinite horizon models (see, e.g., Bertsekas 1976 and Whittle 1982). In the finite horizon case, the method is a backwards sequential solution of Bellman's (1957) equation. Specifically, Equation (4) can be written as follows:

$$(19a) \quad L_t V(\Omega(t)) = R_t(t) + \beta E[V(\Omega(t+1)) | d_t(t) = 1]$$

$$\text{for } t = 1, \dots, T-1, i \in I$$

and

$$(19b) \quad L_T V(\Omega(T)) = R_t(T), i \in I.$$

The solution for $L_1V(\Omega(t))$ is obtained by substituting recursively from T . The existence of a unique optimal policy characterized by a sequence of reservation values $R_i^*(t)$, $i = 1, \dots, T$ has been proved in the literature when $R_i(t)$ follows a general stochastic process (Whittle 1982, Vol. II, Ch. 29). For stationary or nonincreasing processes of $R_i(t)$, the sequence of reservation values is monotonically decreasing.

For the infinite horizon case we seek a time-independent value for $L_1V(\Omega(t))$ for each $i \in I$ that satisfies

$$(20) \quad L_1V(\Omega) = R_i(t) + \beta E[V(\Omega') | d_i = 1], i = 1, \dots, I.$$

It can be proved that under the standard regularity conditions assumed here a steady state solution exists and is unique. This solution is time independent and is consistent with any economic problem in which the researcher assumes that the optimal choice is independent of the particular period of life of the individual.¹⁵

In order to demonstrate the solution method in each case it is useful to simplify the problem. We begin with the two-state search model.

A. Finite Horizon ($T < \infty$) Search Model

In the search model L_1V is the value function when $d_1(t) = 1$, i.e., the individual is not employed, and L_2V is the value function when $d_2(t) = 1$, i.e., the individual is employed.

In the last period, T , the problem is static, i.e., Equation (19b) becomes

$$(21a) \quad L_1V(T) = b,$$

$$(21b) \quad L_2V(T) = w(T).$$

It is clear that $d_2(T) = 1$ if and only if $w(T) \geq b$.

At time $t < T$ Equation (19a) can be written as

$$(22a) \quad L_1V(t) = b + \beta E[V(t+1) | d_1(t) = 1],$$

and

$$(22b) \quad L_2V(t) = w(t) + \beta E[V(t+1) | d_2(t) = 1].$$

Because we know that the problem has an optimal stopping solution, $d_2(t+1) = 1$ if $d_2(t) = 1$; further, $R_2(t+1) = w(t)$ if $d_2(t) = 1$, so that

$$L_2V(t) = w(t) \sum_{s=t}^T \beta^{s-t}.$$

¹⁵ Whittle (1982, Vol. II) provides a formal derivation of the properties of optimal stopping problems for finite and infinite horizon problems.

The solution to the search problem consists of a sequence of reservation wages $w^*(t)$, values of the wage above which a job is accepted and below which it is declined, such that if $w(t) \geq w^*(t)$, $d_2(t) = 1$, $s \geq t$, and if $w(t) < w^*(t)$ $d_1(t) = 1$. The reservation wage at t , $w^*(t)$, solves the equation

$$(23) \quad w^*(t) \left[\sum_{s=t}^T \beta^{s-t} \right] = L_1V(t),$$

which implies that $w^*(T) = b$. For $t < T$ we can sequentially solve for $L_1V(t)$ in terms of $w^*(s)$, $T \geq s > t$. That is, using (22a) we get

$$(24) \quad L_1V(t) = b + \beta E[p \max\{L_1V(t+1), L_2V(t+1)\} + (1-p)L_1V(t+1)]$$

$$= b + \beta p L_1V(t+1) Pr(w(t+1) < w^*(t+1))$$

$$+ \beta p \left(\sum_{s=t+1}^T \beta^{s-t+1} \right) E(w(t+1) | w(t+1))$$

$$\geq w^*(t+1) Pr(w(t+1) \geq w^*(t+1))$$

$$+ \beta(1-p)L_1V(t+1),$$

where

$$Pr(w(t+1) < w^*(t+1)) = \int_{-\infty}^{w^*(t+1)} f(w(t+1)) dw(t+1),$$

$$E(w(t+1) | w(t+1) \geq w^*(t+1))$$

$$= \frac{\int_{w^*(t+1)}^{\infty} w(t+1) f(w(t+1)) dw(t+1)}{Pr(w(t+1) \geq w^*(t+1))}$$

and $L_1V(t+s)$, $s \geq 0$, is given by (22):

Equations (23), (24), and $w^*(T) = b$ jointly determine a closed form algorithm for solving for the reservation wages, $w^*(t)$, $t = 0, 1, \dots, T$. Reservation wages, given the above descriptions, are monotonically decreasing.

B. Infinite Horizon ($T = \infty$) Search Model

For the infinite horizon case there is no terminal solution for the reservation wage. However, it has been proved (see, for example, Bertsekas 1976) that Equation (4) holds for all t and that under certain restrictions on the reward function, there exists a stationary solution for the value function, that is, there exists a V such that $V = V(\Omega(t)) = V(\Omega(t+1))$. For

the search model

$$(25) \quad L_2 V(t) = w(t) \sum_{i=1}^T \beta^{t-i} = \frac{w(t)}{1-\beta}$$

which is equal to $L_2 V(t+1)$ if and only if $w(t) = w(t+1)$. The stationary solution to (20) yields a time independent reservation wage, w^* , such that $d_2(t) = 1$ if and only if $L_2 V(t) \geq w^*/1 - \beta = L_1 V(t)$. To find w^* we substitute this solution into (22a) using (24), i.e.,

$$\begin{aligned} \frac{w^*}{1-\beta} &= b + \beta E \left\{ p \max \left[\frac{w^*}{1-\beta}, \frac{w(t+1)}{1-\beta} \right] + (1-p) \frac{w^*}{1-\beta} \right\} \\ &= b + \beta p \frac{w^*}{1-\beta} Pr(w(t+1) < w^*) \\ &\quad + \beta p \frac{1}{1-\beta} E(w(t+1) | w(t+1) \geq w^*) \\ &\quad \times Pr(w(t+1) \geq w^*) + \beta(1-p) \frac{w^*}{1-\beta}. \end{aligned}$$

Upon simplifying, we get an implicit equation that can be solved for a unique w^* .

$$(26) \quad w^* = b + p \frac{\beta}{1-\beta} \int_{w^*}^{\infty} (w - w^*) f(w) dw.$$

For a given density function it is easy to formulate a numerical algorithm that solves for w^* when β , p , b and the parameters of $f(w)$ are given.¹⁶

C. Finite Horizon With Serial Correlation

Pakes's patent renewal model differs from the finite horizon search model in that the reward is assumed to be serially correlated. In this case the reservation value that characterizes the decision of whether or not to renew in each period depends on the actual realization of the reward in the previous period $R_2(t-1)$. Recall that alternative one signifies non-renewal and alternative two renewal of the patent and that

$$(27a) \quad L_1 V(t) = 0$$

and

$$(27b) \quad L_2 V(t) = R_2(t) + \beta E[\max(L_1 V(t+1), L_2 V(t+1)) | R_2(t), d_2(t) = 1].$$

16. For example, because (25) is a contraction mapping (see Sargent 1987), if we begin at any arbitrary w_0 , and solve for w_1 as the right hand side of (26), then use w_1 as a new value and solve (26) for w_2 , this process is guaranteed to converge to the unique solution.

The optimal stopping rule is characterized by an $R_2^*(t)$ that solves equation $L_1 V(t) = L_2 V(t)$, or

$$(28) \quad 0 = R_2^*(t) + \beta E[\max(0, L_2 V(t+1)) | R_2^*(t), d_2(t) = 1].$$

Pakes proved that under quite general conditions there exists a unique solution to $R_2^*(t)$. He also shows under slightly more restrictive conditions that the reservation value is nondecreasing in t . If $R_2(t) \geq R_2^*(t)$, $L_2 V(t) \geq L_1 V(t)$, then $d_2(t) = 1$.

The problem is to find a closed form solution for the second term in (28). Given the finite horizon, (28) may be solved backwards starting from T . That is, we first find $R_2^*(T-1)$, where $R_2^*(T) = 0$, from the equation

$$R_2^*(T-1) + \beta E[R_2(T) | R_2^*(T-1), d_2(T-1) = 1, R_2(T) \geq 0] \times Pr(R_2(T) \geq 0, R_2^*(T-1)) = 0$$

or

$$(29) \quad R_2^*(T-1) + \beta[(1 - e^{-\alpha R_2^*(T-1)}) \times \int_0^{\alpha R_2^*(T-1)} g(z, T) dz + \int_{\alpha R_2^*(T-1)}^{\infty} z \cdot g(z, T) dz] - c(T) = 0$$

if $R_2^*(T-1) \geq 0$.

Note that because $g(z, t)$ is exponential, the integrals in (29) have closed forms. Similarly, $R_2^*(T-2)$ is calculated from

$$(30) \quad R_2^*(T-2) + \beta E[L_2 V(T-2) | R_2^*(T-2), d_2(T-2) = 1, R_2(T-1) > R_2^*(T-1)] + \beta E[Pr(R_2^*(T-2), R_2(T-1) > R_2^*(T-1))] = 0$$

using the solution to Equation (29) for $R_2^*(T-1)$. This expression is somewhat more complicated to evaluate than is Equation (29). Pakes used a computer software package to solve specifically for the sequence of $R_2^*(t)$, $t = 0, 1, \dots, T$.¹⁷

Except for the serial correlation, Pakes's solution method is the same as that of the finite horizon search model. Eckstein and Wolpin (1987) demonstrate the solution method for the search model when wage offers follow a first-order autoregressive process, i.e., $w(t) = \alpha w(t-1) + u_t$, $0 \leq \alpha < 1$ and u_t has a zero mean density function that is time invariant.

17. The program is called *Maxsyma* and performs analytical differentiation and integration. It should be noted that a simple matching model has $\alpha = 1$, but variances which are finite through time.

Specifically, let $u_i(t)$ be i.i.d. with mean μ_u and variance σ_u^2 . The initial wage draw, $w(0)$, is independent of the u 's and has mean μ_0 and variance σ_0^2 . We assume that wage offers are positive. As already shown, the individual works at T if $w(T) \geq b = w^*(T)$. Serial correlation in wage offers does not alter the optimality stopping property of the model, nor does it affect the existence or uniqueness properties of the optimal reservation wage policy. Therefore, we can solve for the reservation wage sequence recursively as before. If a wage offer is accepted the individual stays with this wage until T . At $T - 1$,

$$V(T - 1) = \max\{w(T - 1)(1 + \beta), b + \beta E(V(T)|w(T - 1))\}.$$

Now

$$\begin{aligned} E(V(T)|w(T - 1)) &= E[\max(b, w(T))|w(T - 1)] \\ &= b \int_{-\infty}^{u^*(T)} dF(u(T)) \\ &\quad + \int_{u^*(T)}^{\infty} (\alpha w(T - 1) + u(T)) dF(u(T)) \end{aligned}$$

where $u^*(T) = w^*(T) - \alpha w(T - 1)$. $F(\cdot)$ is the distribution function of u , and $w(T - 1)$ is given at $T - 1$. Note that the relevant information set includes $w(T - 1)$ because the wage process is of first order. We use the following definition of the conditional expectations of the value function, $E(V(t)|w(t - 1)) = G_{T-1}(w(t - 1))$. Then, $G_{T-1}(w(T - 1)) = E(V(T)|w(T - 1))$ and the reservation wage at time $T - 1$, $w^*(T - 1)$, is the solution to the equation $w^*(T - 1)(1 + \beta) = b + \beta G_{T-1}(w^*(T - 1))$. Now, at $T - 2$, the value function is

$$V(T - 2) = \max\{w(T - 2)(1 + \beta + \beta^2), \beta E(V(T - 1)|w(T - 2))\}.$$

As above, we solve for $w^*(T - 2)$ by first calculating

$$\begin{aligned} E(V(T - 1)|w(T - 2)) &= \int_{-\infty}^{u^*(T-1)} (b + \beta G_{T-1}(\alpha w(T - 2) + u(T - 1)) dF(u(T - 1)) \\ &\quad + \int_{u^*(T-1)}^{\infty} (\alpha w(T - 2) + u(T - 1))(1 + \beta) dF(u(T - 1)) \\ &= G_{T-2}(w(T - 2)). \end{aligned}$$

The reservation wage $w^*(T - 2)$ is found by solving the implicit equation

$$w^*(T - 2)(1 + \beta + \beta^2) = b + \beta G_{T-2}(w^*(T - 2)).$$

Working backwards in this manner leads to a set of reservation wages above which offers are accepted and below which they are declined. As in Pakes (1986) and in the case of time independent wages it can be shown that $G_t(w) > G_{t+1}(w)$ due to the finiteness of the horizon, so that $w^*(t) > w^*(t + 1)$, holding the prior wage realization constant.

D. Rust's (1987) Method

Rust obtained a major simplification by assuming that the error terms are additively separable in the reward function and that they have an extreme value distribution function that is time independent and independent between the different alternatives. This simplification is similar to that obtained in the logit specification of static multivariate discrete choice models (McFadden 1973).

In order to apply Rust's solution method one has to assume that the reward function has the form

$$(31) \quad R_i(t) = u_i(x(t), d(t)) + \epsilon_i(t)$$

where $d(t) = [d_1(t), d_2(t), \dots, d_I(t)]$, that $\epsilon_i(t)$ has the extreme value density function

$$(32) \quad f(\epsilon_i) = \exp\{-e^{-\epsilon_i}\},$$

and that condition (ii), the independence assumption, also holds. The gain from these assumptions is that the second term in Equations (19a) and (20) can be greatly simplified. Given the extreme value distribution assumption, it is straightforward to demonstrate that

$$\begin{aligned} (33) \quad E[V(\Omega(t + 1))|d_i(t) = 1] &= E[\max_{j \in I} \{L_j V(\Omega(t + 1))|d_i(t) = 1\}] \\ &= \sum_{j \in I} (u_j(x(t + 1)) \\ &\quad + \beta EV(\Omega(t + 2)|d_i(t) = 1)) \\ &\quad + \gamma - \ln \Pi_j(t + 1) \Pi_j(t) \end{aligned}$$

where $\Pi_j(t + 1) = Pr(d_j(t + 1) = 1|d_i(t) = 1)$ and where γ is Euler's constant. Note that $\Pi_j(t + 1)$ takes the usual multinomial logit form (see Section IV, Equation (39b), below). In general, the multivariate integration necessary to calculate the left hand side of (33) does not have a closed form solution. However, as is evident in (33) the extreme value distribution assumption obviates the necessity of numerically computing multivariate integrals.

For the finite horizon case the model can be solved backwards from T as described for the search model, using Equation (33) to evaluate the

conditional expectations component of Bellman's equation. Application of the method is straightforward; Berkovec and Stern (1987) have used it in their retirement model.

For the infinite horizon model Rust (1985) developed a nested fixed-point algorithm to solve for the stationary value function. In particular, in the infinite horizon case the value of alternative i , $L_i V(\Omega)$ enters Equation (4) through the term $E[\max(L_i V(\Omega), d_i = 1)]$. The solution for $L_i V(\Omega)$ requires an algorithm that solves for the fixed point of the equation set. The algorithm Rust developed exploits the special assumptions of his model.¹⁹

It should be emphasized that Rust's method requires time independence of the error terms, but it allows any Markovian structure for the observable variables in $x(t)$. The additional assumption that the errors be independent across alternatives, however, is not fundamental, there being a direct analogy to nested logit in static choice models.

E. Gittins Index (Miller, 1984)

A solution method for I discrete alternatives, the Gittins Index, has been developed in Gittins and Jones (1974) and Gittins (1979).²⁰ Gittins Index applies only in the special case where the state that is relevant for alternative i , $x_i(t)$, changes from period t to period $t + 1$ only if $d_i(t) = 1$, that is, $x_i(t + 1) = x_i(t)$ if $d_i(t) = 1$ and $i \neq j$. Changes in the state variables may follow a Markovian transition rule, but this is not necessary for the optimality of the Gittins Index (Varia et al. 1985).

The method of solution is that each alternative i has attached to it an index at each date t that depends only on the state of that alternative. The optimal choice at time t is that alternative with the highest index at time t . The advantage of this method is the fact that a possibly very complicated problem can be reduced to the computation of only I indices at each date, one for each alternative. In the general dynamic programming problem, which alternative is chosen depends on the characteristics of all alternatives. The Gittins Index for each alternative, however, depends only on the characteristic of that alternative.

The structure of Miller's (1984) and Eckstein, Horsky, and Raban's (1988) models allows for the Gittins Index solution. Equation (17) defines the reward function, $R_i(t)$, and Equation (18) specifies the Markovian value function and derives the implied utility function and contraction mapping.

19. The program is available for use on the IBM-PC as part of Gauss software. Rust also developed a backwards solution method. In contrast to the forward method which begins with a specification of the utility function and derives the implied value function as a fixed point to a contraction mapping, the backwards method begins with a specification of the value function and derives the implied utility function and contraction mapping.

20. Later proofs and extensions exist in Whittle (1982) and Varia et al. (1985).

transition rule for the expected value of R_i in the future, which depends only on $M_i(t)$, the number of periods that alternative has been chosen. The Gittins Index is defined as

$$(34) \quad V_i(\gamma_i(t), \delta_i(t)) = \sup_{\tau \geq t} \left[\frac{E \left[\sum_{j=t}^{\tau-1} \beta^{j-t} (R_i(j) - \psi(j)) \mid \Omega(t) \right]}{E \left[\sum_{j=t}^{\tau-1} \beta^{j-t} \mid \Omega(t) \right]} \right]$$

where τ_i is the optimal stopping time for alternative i given that alternative i is chosen. The optimal decision rule is that $d_i(t) = 1$ if $V_i = \max_{j \in I} \{V_j\}$. The index has the intuitive interpretation as the maximum discounted expected reward per unit of expected discounted time, where the maximum is computed at the optimal duration (τ) in an alternative. The dynamic allocation index compares the expected present value of the alternative choices in terms of a normalized return. The choice that currently yields the highest return is undertaken first. Because there is no dependence among the different choices the computation can be done for each choice separately and the value of any given alternative remains the same until that alternative is chosen. Miller (1984) proved that for his model the Gittins index can be written as

$$(35) \quad V_i(\gamma_i(t), \delta_i(t)) = \gamma_i(t) + \delta_i(t) D(\alpha_i + M_i(t), \beta)$$

where $\alpha_i = \sigma_i^2 \delta_i^{-2}$ is called the information factor for alternative i and

$$(36) \quad D(\alpha_i + M_i(t), \beta)$$

$$= \sup_{\tau \geq t} \left[\frac{E \left[\sum_{j=t}^{\tau-1} \beta^{j-t} \delta_j^{-1} (R_i(j) - \psi(j)) \mid \Omega(t) \right]}{E \left[\sum_{j=t}^{\tau-1} \beta^{j-t} \mid \Omega(t) \right]} \right]$$

This decomposition of (34) provides a simplification in the sense that the index D for each alternative is a function of only two parameters, the sum of the information factor of the alternative and duration in the alternative ($\alpha_i + M_i(t)$), and the discount factor (β). $\gamma_i(t)$ and $\delta_i(t)$ can be computed directly from Equation (18), and (36) is found numerically for different values of $\alpha_i + M_i(t)$ and β using a fixed point algorithm described in Miller (1984, Appendix B). Miller uses a cubic spline method as a convenient functional form to compute the value function for the standard index, applying to the contraction mapping algorithm, a method which can be applied to other infinite horizon problems.

IV. Estimation Methods

Most of the following discussion considers as input to the estimation problem a data set consisting of a panel of H individuals or households in which, at a minimum, the choice of each alternative i is observed for each individual for T_h periods.

To establish notation, let the decision set for household h be

$$d^h(t) = [d_1^h(t), d_2^h(t), \dots, d_H^h(t)],$$

and

$$d^h = [d^h(1), d^h(2), \dots, d^h(T_h)]$$

where $d_i^h(t)$ specifies the actual choice of alternative i for individual h at time t . Thus, $d^h(t)$ is the vector defining the choice of alternative at time t for individual h and d^h is the vector describing the choice set over the individual's observed sample period.

A. Maximum Likelihood Estimation

Full information maximum likelihood is used as the method of estimation in most of the studies we have surveyed above. We first describe the general case and then several specific examples. Given data on the actual choices of the H individuals, the likelihood function for these data is given by

$$(37) \quad L = \prod_{h=1}^H Pr(d^h) \\ = \prod_{h=1}^H \{Pr(d^h(T_h)) | d^h(T_h - 1), d(T_h - 2), \dots, d^h(1)\} \\ \quad \{Pr(d^h(T_h - 1)) | d^h(T_h - 2), d(T_h - 3), \dots, d^h(1)\} \\ \quad \dots \{Pr(d^h(2)) | d^h(1)\} Pr(d^h(1)).$$

The second equality is merely a decomposition of the joint probability into conditional and marginal probabilities.

The solution of the dynamic programming model can be used to provide a likelihood value for the parameters of the model conditional on the data d^h , $h = 1, \dots, H$. The parameters of the model which maximize the likelihood function can then be found by a numerical nonlinear optimization algorithm. As in all econometric models identification of the parameters is a crucial issue to which we devote some attention below.

The likelihood function given by Equation (37) already assumes that the decisions of different individuals are independent, i.e., aggregate uncer-

tainty is excluded.²¹ To evaluate the likelihood function without any further simplifications requires in general the evaluation of a multiple integral of order T_h . Even if it is computationally feasible to solve the dynamic programming model it may not be feasible to calculate the likelihood function if T_h is large. Developing a computationally tractable method for evaluating the likelihood function is crucial for the feasibility of empirical work in which structural dynamic discrete choice models are attractive characterizations of behavior.

1. Time Independent Errors

a. Search Model

In the two-state search model the necessary data consist of observations on $d_1^h(t)$, where $d_1^h(t) = 1$ implies that the individual is unemployed. Recall that $d_1^h(t) = 1 - d_2^h(t)$. It is assumed that $d_1^h(t) = 1$ for $t < \tau^h$ and $d_1^h(t) = 0$ for $t \geq \tau^h$, where τ^h is the period in which individual h accepted an offer. Hence, the individual is unemployed for $\tau - 1$ periods. It may be that $\tau > T_h$, in which case the spell is incomplete. Any other sequence of $d_1^h(t)$ has probability zero, i.e., because of its optimal stopping property the model does not admit to a sequence where $d_1^h(t) = 1$ for $t < \tau^h$. As already noted, the solution method discussed in Section III provides a sequence of reservation wages, $w^*(t)$, for the finite horizon case and a single reservation wage w^* for the infinite horizon case. We calculate the probability of being in the two states as follows:

$$(38a) \quad Pr(d_1^h(t) = 1 | d_1^h(t - 1) = 1) = Pr(w(t) \leq w^*(t)) = \int_{-\infty}^{w^*(t)} f(w) dw$$

$$(38b) \quad Pr(d_1^h(t) = 0 | d_1^h(t - 1) = 1) = Pr(w(t) \geq w^*(t)) = \int_{w^*(t)}^{\infty} f(w) dw$$

$$(38c) \quad Pr(d_1^h(t) = 0 | d_1^h(t - 1) = 0) = 1,$$

$$(38d) \quad Pr(d_1^h(t) = 1 | d_1^h(t - 1) = 0) = 0.$$

The conditional probabilities in (38) for $d_1(t)$ are independent of $d_1(t - s)$, $s \geq 2$.

The value of the likelihood function for a given value of the parameters of the model, $\{b, \beta, p, \text{ parameters of the wage offer density function}\}$ is

21. The existence of aggregate shocks is irrelevant to the solution of the dynamic programming models. The only complication arises in the estimation. However, none of the studies we have surveyed incorporate aggregate shocks. See Heckman (1981) for a nonstructural motivation for (37).

found by calculating (38a) and (38b) and substituting the result into (37) according to the actual observation on $d_i^h(t)$, $h = 1, \dots, H$. For the infinite horizon case we get exactly the same conditional probability statements as in (38), but $w^*(t) = w^*$ for all t ; the duration of unemployment has no direct effect on the probability of accepting a job. Estimation proceeds by iterating over the parameter space between the dynamic program to find the reservation wages and the likelihood function until the maximum is achieved.

A simple way of enriching the empirical content of the model is to allow b , p , and $f(w)$ to be functions of observable characteristics including possibly age, schooling, family background, or ability (see, for example, Wolpin 1987). As long as these observables are viewed as exogenous, there is no additional complication in terms of the form of the likelihood function.

b. Labor Force Participation

In the labor force participation model (Eckstein and Wolpin 1988) the likelihood function has exactly the same form as does the search model. Specifically: backward solution for the labor force participation model yields a sequence of T reservation values of $\epsilon(t)$, $\epsilon^*(t)$, $K(t)$, where $K(t)$, experience, takes on the values $t - 1, t - 2, \dots, 0$ for all $t = 1, 2, \dots, T$. Note that $\epsilon^*(t)$, $K(t)$ is a function of $\alpha_1, \alpha_2, \beta_0, \beta_1$ and the parameters of the distribution of $\epsilon(t)$. Assume that the data on each individual consist of 7^* , $K^*(t)$ and $d_i^{7^*}(t)$. Then the likelihood function (37) can be calculated using the following probabilities:

$$(39a) \quad Pr(d_i^{7^*}(t) = 1 | K(t)) = \int_{-\infty}^{\epsilon^*(t), K(t)} f(\epsilon) d\epsilon.$$

$$(39b) \quad Pr(d_i^{7^*}(t) = 0 | K(t)) = \int_{-\infty}^{\epsilon^*(t), K(t)} f(\epsilon) d\epsilon,$$

where $f(\epsilon)$ is the density function.

c. The Dynamic Multinomial Logit (Rust 1987)

We have already pointed out the computational gain of Rust's framework in terms of computing the value function and, hence, in solving the dynamic programming problem. However, the additive nature of the error term in the reward function ($R_i(t)$) together with the assumption that the error follows an extreme value distribution provides a simple analytical form for the conditional probabilities for the choice of alternatives. That is, using Equations (31) and (32) we get

$$(40) \quad Pr(d_i(t) = 1 | x(t)) = \frac{\exp \{u_i(x(t) + \beta E[V(\Omega(t) + 1)] | d_i(t) = 1)\}}{\sum_{j \in I} \exp \{u_j(x(t) + \beta E[V(\Omega(t) + 1)] | d_j(t) = 1)\}}.$$

The term $E[V(\Omega(t) + 1)] | d_i(t) = 1$ is obtained from Equation (33) by solution methods we have already discussed. Note that all the assumptions made by Rust (1987) and that have been discussed above must be satisfied.

2. Serially Correlated Errors

In a search model (or in a patent renewal model) the probability of observing τ periods of unemployment (or τ periods of patent renewal) is equal to

$$Pr(w(0) < w^*(0), w(1) < w^*(1), \dots, w(\tau - 1) < w^*(\tau - 1), w(\tau) > w^*(\tau)).$$

In the serially independent case we have already discussed, this probability is given by the $(\tau - 1)$ independent conditional probabilities shown in (38a) and (38b). If the $w(t)$'s are serially correlated, this expression involves a τ -fold multiple integral.

Pakes (1986), in order to allow for heterogeneity in initial draws, uses simulated frequencies to empirically implement his patent renewal model. For the search model with serially correlated wages this suggestion (see also Lerman and Manski (1981)) amounts to drawing many values of w from its distribution to get a simulated frequency for each possible τ in the data, $\tau = 0, \dots, T$. The simulated frequency for a given τ is simply the count of the number of times the optimum is τ divided by the number of draws. Let $P(\tau, \theta)$ be the simulated frequency, given the parameter vector θ , for τ unemployment periods. For each τ there are $n(\tau)$ observations in the data. Then the pseudo likelihood function is given by

$$(41) \quad L(\theta) = \prod_{\tau=1}^T P(\tau, \theta)^{n(\tau)}.$$

Pakes proved the asymptotic consistency and normality of the estimated parameters that maximize the pseudo likelihood function.

In Miller's (1984) occupational choice model, wages follow a modified random walk. Given a duration of τ periods on a job, the hazard rate for switching to a new job in the same occupation is defined as

$$(42) \quad h_i(\tau) = Pr[y_i(t) + \delta_i(t)D(\alpha_i + M_i(t), \beta) < \gamma_i + \delta_i D(\alpha_i, \beta) | M_i(t) = \tau].$$

The term $\gamma_i + \delta_i D(\alpha, \beta)$ is the value of a new job in the occupation, i.e., $M_i(t) = 0$. Letting $P(\tau)$ be the unconditional probability of observing a duration of τ periods,

$$(43) \quad P(\tau) = h(\tau) \prod_{j=1}^{\tau-1} (1 - h(j)).$$

If $n(\tau)$ is the number of individuals staying τ periods on a job, then the likelihood function is given by

$$(44) \quad L = \prod_{\tau=1}^{\infty} P(\tau)^{n(\tau)}.$$

The main difficulty is in computing $h(\tau)$ for different levels of τ . To do that the dynamic index, $\gamma_i(t)$ and $\delta_i(t)$ have to be computed jointly with the implied probability of $h(\tau)$. Miller developed an algorithm for this purpose and computed the index and $h(\tau)$ for different values of the underlying parameters. Miller carried out the estimation by searching for the highest value of the likelihood function over a prespecified grid of values of the parameters α and β .

B. Identification

Identification is defined in terms of the conditions for existence of consistent estimators for a given model and estimation method. Hence, the conditions for identification of nonlinear models using the maximum likelihood estimation method is equivalent to the conditions for maximum likelihood to yield consistent estimates (see, e.g., Amemiya 1985). There is no general method to check the conditions prior to estimation. However, it is possible in some cases to determine whether some of the parameters are not identified. There may be situations where parameter constraints due to the theory imply that the same likelihood value is generated at every point in parts of the parameter space. For example, as in any econometric model, if the reservation wage in the search model were a function of a ratio or sum of parameters, then only the ratio or sum would be identified. This is in fact the case in the labor force participation model where α_1 and α_2 enter only in differenced form. In some models the reservation values have a different form at different dates in terms of the underlying parameters (see, for example, Wolpin 1987). Hence, identification would depend in this case on whether the data contain information on more than one period. To be more specific consider the search or the labor force participation model. As specified in the likelihood function, we require that all of the parameters, including the wage distribution parameters, be identified from the data on d_i 's only. It is easy to demon-

strate that some parameters of the wage function (e.g., β_0) cannot be identified unless we include observations on wages. If the wage is observed and included in the estimation, then the likelihood function must be modified because the wage is assumed to be random. If we observed all of the wage offers regardless of whether they were accepted we could estimate the wage function parameters directly. Observing only accepted wages, as is usual, leads to standard selectivity issues, but wage parameters can be identified either parametrically if a wage distribution is assumed (Heckman 1979) or nonparametrically (Ichimura and Lee 1989; Heckman and Honoré 1988). However, efficient estimation requires that we estimate the model jointly because the parameters of the wage function directly affect the acceptance decision.

Using wage information, the likelihood function is the joint density of the accepted wage and the d_i 's. Now, the maximization of the joint likelihood must take into account the fact that the reservation wage in each period must be less than the smallest observed accepted wage in each period. (If rejected wages are observed then the reservation wage must also be greater than the largest rejected wage offer in each period). Because there is only one source of randomness, a single outlier wage observation can have a disproportionate effect on the estimated parameters. This problem may be avoided by introducing an additional source of randomness. The least complicated approach, and one that is not implausible, is to assume that wages are measured with error (Stern 1989 forthcoming; Wolpin 1987). It is the least complicated because it doesn't affect the computation of the reservation wage. Alternatively, one can include an error term representing preference heterogeneity as in Rust's specification or in Wolpin's fertility model.

In Rust's (1987) model the state vector $x(t)$ is assumed to be observed at each date and governed by a parametric Markovian transition rule as in Equation (8). Identification of that process requires data on $x(t)$. The joint likelihood function for the observed decision $d_i(t)$ and the observed $x(t)$ is a simple extension of Equation (37) explicitly taking into account the evolution of the state space as given by (8).

C. Method of Simulated Moments

In several of the models the calculation of the likelihood function requires multivariate integration. This is the case, for example, when the error is serially correlated as in Pakes. These calculations are computationally demanding and estimation in many interesting models may not be feasible. In several recent papers, McFadden (1988) and Pakes and Pollard (1988) have developed an estimation method based on simulated, rather than computed, moments. Hansen (1982) provides a general discussion of

