

Rota, Paola

Working Paper

What do we gain by being discrete? An introduction to the econometrics of discrete decision processes

Nota di Lavoro, No. 22.2001

Provided in Cooperation with:

Fondazione Eni Enrico Mattei (FEEM)

Suggested Citation: Rota, Paola (2001) : What do we gain by being discrete? An introduction to the econometrics of discrete decision processes, Nota di Lavoro, No. 22.2001, Fondazione Eni Enrico Mattei (FEEM), Milano

This Version is available at:

<https://hdl.handle.net/10419/155185>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Fondazione Eni Enrico Mattei

**What Do We Gain by
Being Discrete?
An Introduction to the
Econometrics of Discrete
Decision Processes**

Paola Rota*

NOTA DI LAVORO 22.2001

APRIL 2001

ETA – Economic Theory and Applications
--

*University of Modena and Reggio Emilia, Italy

This paper can be downloaded without charge at:

The Fondazione Eni Enrico Mattei Note di Lavoro Series Index:

http://www.feem.it/web/attiv/_attiv.html

Social Science Research Network Electronic Paper Collection:

http://papers.ssrn.com/paper.taf?abstract_id

Fondazione Eni Enrico Mattei
Corso Magenta, 63, 20123 Milano, tel. +39/02/52036934 – fax +39/02/52036946
E-mail: letter@feem.it
C.F. 97080600154

What Do We Gain by Being Discrete? An Introduction to the Econometrics of Discrete Decision Processes

Paola Rota

Universita' di Modena e Reggio Emilia

November 2000

Abstract

In this paper we analyze methods which allow us to estimate and identify the sources of censoring in dynamic models. We explicitly take into account the existence of corner solutions by considering a discrete-time-discrete-choice dynamic structural model. The availability of microeconomic datasets allows us to focus on decisions at the individual level and directly exploit the information contained in the corner solutions. We show how a discrete decision process (DDP) represents a natural framework within which to analyze agents' behaviour when optimal inaction generates censoring in observed decisions. A discrete decision process is characterized by a control variable which only takes a finite number of values. Some problems are naturally discrete, such as the optimal engine replacement or job the search problem in which the individual decides whether or not to accept a job offer. Other problems may be described very efficiently by a discrete decision problem. This is clear in the case of fixed costs of adjusting inputs which imply the discrete decision of whether or not to vary the production factor.

Keywords: discrete choice, dynamic models, transition probabilities

JEL classification: C35, C51

This paper is the text of two lectures delivered at the Universita' di Bergamo, 6-7 November 2000. I would like to thank Elena Cefis, Annalisa Cristini, Alessandra Del Boca and Marzio Galeotti for comments. I gratefully acknowledge financial support from the Ministero dell'Universita' e della Ricerca Scientifica e Tecnologica: MURST project no. 9913467782.

1 Introduction

The recent economic literature has paid increasing attention to modelling agents' infrequent reactions in response to shocks. This type of behaviour is typically associated with the presence of lump-sum and/or kinked adjustment costs, partial or total irreversibilities, menu costs, indivisibilities and non-negativity constraints. Infrequent behaviour is apparent in longitudinal datasets which report firms (or other agents) decisions over time. For instance, there is evidence that firms may not respond to observed changes in demand and/or costs, even when these are of a large magnitude. A consequence of inaction is the existence of corner solutions which cause a break down in the standard marginal conditions for optimality given by Euler equations. This implies censoring in the observed decision variables. Censoring is simply the result of the possibility that an agent may find that the optimal choice is "doing nothing" and postpone economic decisions to the future. Dixit and Pindyck (1994) describe, in the case of investment, a firm with an opportunity to invest as holding an option. When a firm makes an irreversible investment expenditure, it exercises its option to invest. It gives up the possibility of waiting for the arrival of new information which might affect the desirability or timing of the expenditure; it cannot disinvest should market conditions change adversely. This lost option value is an opportunity cost that should be included in the cost of investment.

In this paper we analyze methods which allow us to estimate and identify the sources of censoring in dynamic models. We explicitly take into account the existence of corner solutions by considering a discrete-time-discrete-choice dynamic structural model. The availability of microeconomic datasets allows us to focus on decisions at the individual level and directly exploit the information contained in the corner solutions. While discrete time is obviously a characteristic of the data, discrete choice also represents the natural framework within which to analyze agents decisions of remaining inactive in certain periods even in the presence of shocks. Some models may also comprise both continuous and discrete choices, see Pakes (1994).

Structural models attempt to derive behavioural predictions from an explicit solution to an underlying optimization problem. Given specific functional forms for the primitives, the agents preferences and beliefs, we can estimate and identify the unknown parameters. These restrictions impose a tight and interpretable structure on our models although at the cost that they may omit other ancillary features of the data. We, thus, may use structural

models to implement counterfactual experiments and evaluate the effects of changes in the behavioural parameters. (See Keane and Wolpin (1997) for a brief discussion of structural estimation in applied microeconomics.)

The foundations of the structural analysis of discrete choices are to be found in the work of the most recent Nobel laureate D. McFadden. We anticipate McFadden approach which we will discuss in Sections 3 and 4. We start from agents' decisions as recorded in the data. Both decisions and information relevant for describing a certain economic phenomenon are observed with some error. This implies that we can only define observed decisions in probabilistic terms. However, assuming rationality, agents' decisions must be consistent with some individual optimizing behaviour. McFadden proposes, in a static context, an axiomatic approach (a *stochastic revealed preference* axiom) which implies the consistency between probabilistic choices and a model of optimizing behaviour. Although he considers a static framework, his methodology may be extended to embed dynamics. We start to consider dynamics in Sections 5 and 6.

In Sections 6 and 7 we define DDP models: their distinct characteristic is given by a control variable which takes a finite number of alternatives, for example a firm decides whether to upgrade or downgrade or maintain unaltered the stock of capital. Another crucial feature of DDP models lies in the specification of uncertainty. There are two sources of uncertainty. The first, is common to all models which assume uncertainty, is given by the possibility of exogenous shocks to the variables relevant for the decision process. This may entail a revision of the optimal plans after the realization of a shock. The second source of uncertainty is typical of the discrete choice methodology, and it arises from the econometrician's lack of knowledge on the state variables relevant for the decision process. This may be the consequence of difficulties in obtaining exact information about some state variables: prohibitive costs, inaccessibility and imperfect monitoring, and it introduces an observation noise which plays a crucial role in the identification and estimation of the structural parameters. While agents know the state variables relevant for their optimal plans, the econometrician only has imperfect knowledge about these variables. In Sections 8, 9 and 10 we discuss simplifying assumptions which allow us both mathematical and econometric tractability, such as Rust's *Conditional Independence*, which allow us to obtain a dynamic version of the basic static discrete choice model proposed by McFadden.

In Sections 11 and 12 we show how, in the context of McFadden's ax-

iomatic approach, an optimal policy exists and it is the solution to the DDP. Together with the simplifying assumptions just mentioned we will be able to obtain a closed form solution, which allow us to directly estimate the structural parameters in a two stage estimation procedure, as shown in Section 13. In Section 14 we discuss an application of the DDP methodology and in Section 15 we draw some conclusions. Although, there has been number of competing methods for estimation of DDP, all based on applications of maximum likelihood (ML) estimation, we concentrate on the approach proposed by Hotz and Miller (1993) which avoids the computational burden of directly estimating the structural model implied by the ML estimator.¹

In the next section we briefly discuss some models in which agents' behaviour implies a discrete choice set.

2 Models with Discontinuous Behaviour: a brief survey

An increasing number of studies emphasize the role of inaction in agents optimal decisions. Models of discontinuous behaviour has been flourishing in recent years in many fields. Given their large number, in this section we will only mention a small set. We divide the literature into two groups: Models with discontinuous behaviour, such as (S,s) rules, which would be best modelled as DDP, but which are not necessarily considered within this structure in the studies reported below and models which directly on DDP and propose methods for their estimation.

Recent studies on dynamic labour demand emphasize the importance of non-convex components in the structure of hiring and firing costs in the form of either fixed or kinked adjustment costs (Hamermesh, 1989, Bentolila and Bertola, 1990, Hopenhayn and Rogerson, 1993, Aguirregabiria and Alonso-Borrego, 1999, Rota (2000) and, for a discussion of the literature, Hamermesh and Pfann, 1995). These studies have emphasized how non-convexities in the structure of adjustment costs may have been substantially enhanced by labour market regulations.² Some hiring and firing provisions add an

¹Other methods which avoid direct solution of the structural model are Manski's (1993) path utility approach, Rusts's (1988) backward estimation and Hotz, Miller Sanders and Smith's (1992) simulated value function.

²Del Boca and Rota (1998) provides an account of the legislation on hiring and firing in Italy.

element of fixity to the cost structure, providing an incentive for firms to react intermittently to changes in the exogenous variables.

In the case of a fixed cost, employment change tends to be concentrated in a single-period, so that firms avoid paying this cost too frequently. Moreover, firms only make those changes in the labor input which are justified by sufficiently large departures of desired employment from their most recent choice of the number of employees. The adjustment process is lumpy and intermittent: in the face of a shock, a firm may decide that it is optimal to maintain the same number of employees and to postpone adjustment to the future. Continuous and gradual reactions to exogenous shock would penalize the firm since each small change would imply the payment of fixed costs.

Kinked linear costs imply a constant marginal cost of adjustment. Hence the adjustment is not lumpy, although it stops before the level of employment which would be attained in the absence of those costs is reached. Also in this case, there is a band of inaction around the no-costs target also in this case, generated by two target levels: one for hiring and one for firing. For levels of employment close to the target levels it will still be optimal for a firm not to adjust.

In either case, the standard marginal conditions for optimality given by the Euler equation fail to hold. The adjustment path is characterized by periods of inaction and discrete jumps is often defined as (S,s) rule. Consider the case of demand for inputs, say labour, in the presence of fixed hiring and firing costs. A two-sided (S,s) rule may be defined as the following: if the number of employees is above (below) or equal to a critical threshold s^D (s^U) then the firm decides to reduce (increase) employment to its desired level S, otherwise it leaves it unchanged - superscripts U and D indicate upward and downward adjustment respectively. Hence, there is a zone of non-adjustment delimited by the two critical values s^D and s^U .³

A large part of the debate on (S,s) rules has concentrated on expenditure on durable goods and inventory management. These studies emphasize the lumpy nature of durable goods purchases: individuals update their durable stocks infrequently and when they do make purchases their purchases are large. In other words, whenever an individual's existing stock reaches a critical level s, she will make a purchase and fully update her stock to a level S. If the existing stock remains above this threshold then no purchase occurs.

³In Rota (2000) the conditions under which we may empirically represent firms' intertemporal employment decisions as an (S,s) rule are defined.

See Lam (1991) for the case of durable goods expenditure, Eberly (1994) and Attanasio (1997) for the analysis of households' decisions to update their holdings of automobile, and Aguirregabiria (1999) who combines an (S,s) inventory model and menu costs.⁴

Recently consideration of irreversibility in the analysis of investment behaviour has posited the attention on the possibility of firm's inaction and intermittent capital adjustment. What causes partial or total irreversibility: capital goods be, at least partially, firm-specific and, hence, the cost of capital may be partially or totally sunk or the price at which capital goods are purchased is different (higher) from the price at which those goods can be sold. In these cases, investment occurs infrequently and the standard formulation of q theory breaks down. Various studies provide evidence that a large portion of investment at the plant level is concentrated in a few episodes. Cooper, Haltiwanger and Power (1995) show that at the micro level the probability of a plant experiencing a large investment episode is increasing in the time since its previous episode. At plant level, bursts of investment are then followed, on average, by periods of low investment. This is in contrast with the usual presumption of positive serial correlation in investment activities which emerges from the standard convex adjustment cost model. Nilsen and Schiantarelli (1996) find evidence of lumpiness both at plant and firm-level. Barnett and Sakellaris (1998) find that investment has a non-linear relation to q , Abel and Eberly (1999) allow for irreversibility and fixed costs in a model of heterogeneous capital goods and Caballero and Engel (1993 and 1994) show that lumpy adjustments play an important role in firms' investment behaviour.⁵

There are numerous models in which the decision problem is inherently discrete such as the optimal engine replacement model (Rust, 1987).⁶ Harold Zurcher is the maintenance manager of the Madison Metropolitan Bus Com-

⁴A number of studies has proved the optimality of (S,s) rules: Scarf (1960), Constantinides and Richard (1978), Grossman and Laroque, (1990) and Bar-Ilan and Blinder (1992).

⁵An important issue, on which recent studies on intermittent behaviour have concentrated is the possibility that microeconomic frictions influence aggregate dynamics. The effect of heterogeneity of agents, idiosyncratic uncertainty and lack of coordination can imply highly nonlinear behavior in aggregate time series. Hence, reference to microeconomic models becomes relevant if aggregate dynamics are to be correctly understood, see Caballero and Engel (1993) and Bertola and Caballero (1994) and Caballero (1997).

⁶See the Eckstein and Wolpin (1989) and Rust (1994) for a survey and a more detailed discussion of these models.

pany. He has to decide how long to operate a bus before replacing its engine with a new completely overhauled bus engine. The estimated model implies that the expected monthly maintenance cost increase by \$1.87 for every additional 5000 miles. A finding very close to the Zurcher's perceptions of operating costs.

Other studies have been carried out on problems concerning airline's decisions on whether or not to remove and overhaul an aircraft engine (Kennet, 1992); whether or not to refuel a nuclear reactor (Sturm, 1991); whether or not to renew a patent and pay the renewal fee (Pakes, 1986); whether or not to dig an irrigation well in semiarid India (Fafchamps and Pender, 1997). Discrete decision process are a very natural representation of job search problem, Wolpin (1987), in which the individual decides whether or not to accept a job offer when received, and job matching models Miller (1984).

Hotz and Miller (1991) estimate a DDP model involving a family contraceptive choice. The family faces three choices: one extreme choice, such as sterilization, which is regarded as irreversible and is modelled as an absorbing state, a choice of no contraception and a choice of temporary contraception. This paper is very important because in that Hotz and Miller develop a new estimation method which reduces the computational burden typical of these models, which otherwise would require repeated solution of the optimization problem at each trial parameter value. We will refer to their estimation methodology in the discussion on the identification and estimation of DDP models.

3 Static Models of Discrete Choice

Although our focus is on intertemporal models, it will be very useful to illustrate briefly a class of static models which give us an important insight in the understanding of economic processes which involve discrete decisions. We follow the work by McFadden (1973,1974) in the context of probabilistic consumer theory of discrete choice. Assume that consumers are rational, ie that they make choices that maximize their perceived utility subject to their budget constraints. Consumers know the variables which are relevant for their optimization problem, although they may be uncertain about future values of those variables and their decisions may be subject to exogenous shocks.

However, there is an additional and very important source of error which

is typical of micro datasets which record information on a large number of individuals, often over a number of periods. Those data typically have recording errors of various type, which implies that data observed by the econometrician will contain an error term; in addition to this source of error, there may be unobserved state variables and/or unobserved characteristics which are important in modelling individuals' behaviour, such as fixed effects, which result in the econometrician having less than full information on all the relevant variables. These error components make utility random from the standpoint of the econometrician, even in the case in which agents take decisions in a situation of perfect information.

Thus, the decisions we observe from data, should be described in terms of probabilistic choice models. However, assuming rationality, these decisions must be consistent with some individual optimizing behaviour. Is a model of random utility maximization indentifiable from the observed distributions of demands? McFadden and Richter (1970) propose a *stochastic revealed preference* axiom which implies the consistency between a probabilistic choice system and a model of optimizing behaviour. In other words, the demand by a certain individual, observable in the data, is consistent with that individual's utility maximization, and it represents the most likely outcome, given her choice set and the attributes observable by the researcher.

Although the literature is based on random utility maximization models, in order to be maintain a consistent notation in what follows, we consider a model of profit maximizing decisions. In Appendix 1 we further discuss McFadden's model of probabilistic consumer theory.

Consider a dataset of firms faced with J choices. The profit from choice $d=j$, Π_{ij} may be expressed as:

$$\Pi_{ij} = \pi_{ij} + \varepsilon_{ij} \quad (1)$$

where $i, i=1, \dots, N$, indexes firms and $j, j=1, \dots, J$. The profit from choice $d=j$ is partitioned in two components: an observable term, π_{ij} , and a random component, ε_{ij} . Profit maximization implies that

$$P(d_i = j) = \Pr(\Pi_{ij} > \Pi_{im}) \quad \text{for all other } m \neq j \quad (2)$$

where m indicates a decision alternative to j . The probability of observing decision j is defined once we make distributional assumptions on the error term, ε . McFadden (1973) showed that if the J disturbances ε_{ij} are iid and

follow an extreme value, or Weibull, distribution,⁷

$$F(\varepsilon_{ij}) = \exp(-\exp(-\varepsilon_{ij})) \quad (3)$$

then P_{ij} is a logit. This implies that if the firm makes a choice, $d=j$, then the pay-off from that choice Π_{ij} must be the largest of all J pay-offs. The probability of choice j , may be thus expressed as

$$P(d_i = j) = \frac{\exp \pi_{ij}}{\sum_j \exp \pi_{ij}} = \frac{\exp(\beta' x_{ij})}{\sum_j \exp(\beta' x_{ij})} \quad (4)$$

x_{ij} is a vector of individual and choice specific attributes. Suppose that the regressor vector x_{ij} has dimension $(k \times 1)$, then there are only $(J-1) \times k$ identifiable parameters. The logit model has a very useful property, ie the probability of any pair of decisions (j, m) is equal to the log-odds ratio:

$$Q_i(j, m) = \log\left(\frac{P_{ij}}{P_{im}}\right) = x'_i(\beta^j - \beta^m) = \pi_{ij} - \pi_{im} \quad (5)$$

The logit model has the benefit of having a simple but restrictive structure. Since we are assuming iid disturbances, the log-odds ratio does not depend on the total number of choices considered, J . If the individual were offered an expanded choice set, equation 5 would not change. This property is referred to as "*independence of irrelevant alternatives*", Luce (1959); see also Manski (1981, 1984). The standard example is that of the red and blue buses, attributed to McFadden but noted by Debreu (1960). We follow Cramer (1991) for an illustration. Consider the choice among several modes of transport, with j denoting some form of private transport such as driving or cycling, and m travel by a red bus, the only mode of public transport. Suppose that a new service is introduced: a blue bus, identical to the existing buses except for the colour. The choice set is now wider, but the log-odds ratio formula for (j, m) is unchanged. In other words, the new service is suppose to proportionately reduce the probabilities of all existing transport

⁷A standard type I extreme value distribution, also known as Weibull or Gumbel, has the form:

$$F(x) = \exp(-\exp(-(x - \mu)/\lambda))$$

where μ is a location parameter and λ is a scale parameter. In the standard form $\mu = 0$ and $\lambda = 1$.

modes. But it is much more likely that the blue bus will affect the red bus utilization more than that of the other means of transport.

This structural interpretation of the static logit model represents the foundation of the analysis of discrete choices over time and, and as we shall see, equations 4 and 5 provide the large part of the theoretical setup we will employ. However, within a dynamic framework we need to take into account that a decision today affects choices in the future and this complicates the analysis.

4 Social Surplus

The choice probabilities $P(d | x)$, ie the probability of observing choice d given the information in the data available to the econometrician, are given by the gradient of a Social Surplus Function G proposed by McFadden (1974). The Social Surplus was used to prove the consistency of the conditional choice probabilities as a representation of a random utility maximization problem. Indicate the distribution of the error terms as $q(\varepsilon | x)$ - which is a logit if the ε 's follow an extreme value distribution, the following function G constitutes a Social Surplus:

$$G(\pi_j(x) | x) = \int \max_j [\pi_j(x)] q(\varepsilon_j | x) \quad (6)$$

G satisfies the following properties:

1. It is a convex function of $v_j(x)$ and d
2. It has the additive property that $G(\{\pi_j(x) + \alpha\} | x) = G(\pi_j(x) | x) + \alpha$, where α is a scalar.
3. The partial derivative of G with respect to $\pi_j(x)$ exists and equals the conditional choice probability:

$$\frac{\partial G(\pi_j(x) | x)}{\partial \pi_j(x)} = P(d | x) \quad (7)$$

Assuming a population of random utility maximizers indexed by ε , $G(\pi_j(x) | x)$ may be thought of as the expected indirect utility function from choosing alternative j .⁸

⁸Also stated to be the discrete choice version of the Roy's identity, see McFadden (1981)

5 A Dynamic Model of Discrete Choice

When the problem becomes dynamic the consequences of choosing an action from the set of available alternatives is twofold: the firm receives an immediate reward, current profits, but also it specifies a probability distribution for the subsequent system state. Then, in order to characterize the profitability of a current choice we must assess the expected profitability of subsequent choices. Hence, in each period and in any state, the decision rule has the property that it must also be optimal for the continuation process, treating the current state as starting point. Decisions are taken and revised in stages, based on information as it becomes available and agents' beliefs about the future.

The major tool for dealing with sequential optimization under uncertainty is dynamic programming or (DP) algorithm, ie a method which reduces a multidimensional problem to a recursive solution of a sequence of two-period problems (Bellman's Principle of Optimality, 1957). The DP algorithm allows us to compute the optimal decision rule by backward induction, starting at a finite terminal point, T. The logic of backward induction may be extended to infinite horizon problems. The DP algorithm may be applied to either continuous, or discrete and mixed discrete continuous choice sets.

Assume discrete time and infinite horizon framework and consider the following specification:

$$\max_{d_t} E_t \left\{ \sum_t \beta^t \Pi(s_t, d_t) \right\} \quad (8)$$

where Π indicates profits, s_t is a vector of state variables, d_t is the decision variable, β is the discount factor, $\beta \in (0, 1)$ and the expectation E is taken with respect to information available at time t . We omit the firm index in order to simplify notation.⁹ We may express the sequence problem 8 by the following value function (Bellman, 1957):

$$V(s_t) = \max_{d_t} \{ \Pi(s_t, d_t) + \beta E[V(s_{t+1}) | s_t, d_t] \} \quad (9)$$

⁹The firm selects an action d_t at time t with knowledge of the state s_t ; a decision rule $\delta(s_t)$ is a mapping of the state vector s_t into the decision d_t , representing a decision rule for choosing an action at each period t , for each possible value of state variables that may occur; a policy $\delta(\delta_1, \dots, \delta_{t-i}, \dots)$ is a sequence of δ^t s and represents a decision rule for choosing at each period t an action for each possible value of the state variables over the entire horizon.

where the first component on the right hand side indicates current profit and the second part represents the continuation value. The valuation function enables us to decentralize a complicated multiperiod problem into a sequence of two-period decision problems, providing the correct valuation of the future consequences of current actions. Ideally we would like to use the dynamic programming algorithm to determine closed-form solutions for $V(s_t)$ but in many cases one has to resort to numerical solutions and this may be time consuming since the maximization must be carried out for each value of s_t . Indeed, for complex multidimensional problems the computational burden may be prohibitive. In what follows we try to reduce this dynamic model to obtain an estimable form.

6 A Markovian Model

Equation 9 has a Markov structure. Markov processes have the defining property that all the information relevant to the determination of the probability distribution of future values is summarized in the current state. This implies that the firm does not need to remember the entire previous history to solve its optimization problem, but only a summary statistics belonging to a finite vector space, $s_t \in S$, ie the value of the set of state variables s at time t , and its law of motion, ie the probability distribution which characterized how the state changes from period to period:

$$p(s' = s_{t+1} \mid s = s_t, d_t) \quad (10)$$

where the prime indicates variables not known in the current period. Probability 10 stochastically determines next period's state s' as a function of the current state vector s_t and the current decision, d_t . The Markovian property is clearly very desirable since it effectively summarizes past information that is relevant for future optimization making previous history irrelevant. It allows us to consider time in two blocks: the present and the future.¹⁰ In

¹⁰The state variable s_t , is sufficient, ie it effectively summarizes past information that is relevant for future optimization making previous history irrelevant. Given the value of s_t one may calculate the optimal decision and also the future value s_{t+1} . It represents a very economical description of the state of the system making previous history irrelevant as far as future profits are concerned. Otherwise, the optimal decision would be a function of the previous history of the system determined by all the values of previous decisions and states. Previous history is itself sufficient too but it takes values in an increasing space

equation 10 we are assuming stationarity. According to this assumption, the future looks the same whether the firm is in state s_t at time t or in state s_{t+k} at time $t+k$, provided that $s_t = s_{t+k}$, ie the current state is the only variable which affects the firm's perception of the future. Then we may write $s' = s_{t+1}$, $s = s_t$ and $d = d_t$. Hence the decision rule and the value function are time invariant.¹¹

7 A Discrete Decision Process (DDP)

In this section we introduce a particular class of Markov models. Suppose that the decision or control variable, d , is restricted to a countable set of alternatives. For example, the firm has to take one of two mutually exclusive choices: decide whether to invest, knowing that after purchasing a machinery, costs will be partially or completely sunk, or not invest and keep the option to invest for the future. Another example is the decision the firm has to hire, fire or not to vary the stock of employees, in the case of non-convex adjustment costs. In this case d may assume three alternative values: hire, not to vary the number of employees, fire. If the control variable only takes a countable set of values, then the Markov structure specializes to a discrete decision process (DDP). Thus we obtain the following Bellman's equation:¹²

$$V(s) = \max_d \left\{ \Pi(s, d) + \beta \int V(s') p(s' | s, d) \right\} \quad (11)$$

where d is equal to a finite (typically small) number of choices.

A discrete-choice dynamic structural framework such as 11 represents a natural framework within which to analyze firms' behaviour when optimal

as t increases, while s_t takes values in a fixed space. It clearly simplifies the problem to use variables which take values in a fixed low dimensional space, and indeed it is desirable that this dimension be as low as possible.

¹¹The Markov structure may seem very restrictive. Indeed, suppose lagged values of the state variables were relevant, say s_{t-1} , we can expand the state vector to $s = (s_t, s_{t-1})'$ and still maintain the Markovian property. However, they may problems with high-dimensional state vectors, which make computation very burdensome.

¹²The infinite horizon case introduces complications since an infinite number of periods implies an infinite number of decisions. This requires the study of the convergence property of the DP algorithm. The Contraction Mapping Fixed Point theorem proves that the function $V(s)$ is the unique solution to Bellman's equation. See Bertsekas, (1987) for a proof.

inaction generates censoring in observed decisions. We mentioned in the introduction that the optimizing response in the presence of fixed adjustment costs may be inaction and thus there is the possibility of corner solutions in the demand for inputs. In this case, the standard marginal conditions for optimality given by the Euler equation fail to hold. Another example is given by (S,s) rules in the case of expenditure on durable goods and inventory behaviour. A large number of studies emphasize the lumpy nature of durable goods purchases and inventory management. In the case of durable goods, individuals update their durable stocks infrequently and when they do make purchases their purchases are large. Infrequent behaviour is apparent in longitudinal datasets which reports firms' (or other agents) decisions over time and show that firms may not respond to observed changes in demand and/or costs, even when these are of a large magnitude. The DDP framework allows us to model directly inaction and to exploit the information contained in the discrete choice. See Eckstein and Wolpin (1989), for a survey and Rust (1994) for a discussion.

8 Unobserved State Variables

The econometrician's goal is to estimate the unknown parameter vector $\theta \in \Theta$, in a model that is fully specified up to Θ , a compact convex subset of a Euclidean vector space. As in the case of the static model discussed in Section 3, if we were able to observe all components of s (given forecasts of y) and if the primitives were correctly specified, the estimated decision rule would allow us to perfectly predict firms' behaviour at each point in time and to know the precise consequences of their decisions. This is a consequence of Blackwell's theorem (Blackwell, 1962 and 1965) which implies that the optimal decision rule is a deterministic function of the agent's state vector, s .¹³ However, datasets are not rich enough to provide full information on the state variables relevant for the firm's decision making and real data usually show different values of the decision, d , associated with the same

¹³Under certain regularity conditions, such as boundedness of π , the Bellman's equation is mathematically equivalent to a fixed point contraction mapping. This implies existence and uniqueness of the DP algorithm. Blackwell's theorem states that the stationary decision rule computed from Bellman's equation is an optimal decision rule. A stationary policy is Blackwell optimal if it is simultaneously optimal for all the β -discounted problems, with $\beta \in (0, 1)$. See Bertsekas (1987) and Stokey and Lucas (1989) for a discussion.

values of the state variables.¹⁴ The typical solution is to add the error term such as unobserved state variables and/or measurement errors, in order to reconcile the difference between the observed behaviour and the prediction of the discrete decision process. From our standpoint, we assume that this error term is a consequence of the fact that our data fail to fully measure all the characteristics of a firm and partition the vector of the state variables as the following:

$$s = [x, \varepsilon] \quad (12)$$

In expression 12 we allow the state vector, s , to comprise an observable component x , known to the econometrician (and to the firm), which include all variables, knowledge of which can help to determine a certain decision, and an unobservable term ε , which reflects the econometrician's imperfect knowledge about the state variables relevant for the firm's decision process. Hence, the decision rule is deterministic from the standpoint of the firm but stochastic from our standpoint. This disturbance term has the same role as in the case of the static models discussed in Section 3.

Following the discussion in Section 3, we make the following assumption:

Assumption 1: *The profit function is additively separable and may be written as*

$$\Pi(s, d) = \pi(x, d) + \varepsilon(d) \quad (13)$$

where the unobserved state variable ε is a vector with at least as many components as the number of alternative choices (McFadden, 1981; Rust, 1987) and where $E[\Pi(x, d)] \equiv \pi(x, d)$ is the one-period profit, conditional on the observables x and the choice d .

Because of the presence of the random disturbance term ε , profits are a random variable and we have thus to consider their expected value, where the expectation, $E(\cdot)$, is taken with respect to the joint distribution of the random variables. As we have already seen, additive separability between observables and unobservables was introduced by McFadden (1973) and (1981) in order to

¹⁴Manski (1977) lists a number of reasons for incorporating this random component: it reflects unobserved characteristics, unobserved taste variation and similar imperfections which force the analyst to treat the choice process as random. See also McFadden (1973, 1981) in the context of static structural discrete choice models and Bertsekas (1987) for a discussion of problems with imperfect state information.

define the random preference maximization model, in which preferences are influenced by a unobserved variables. It is clearly a strong assumption, but relaxing it requires one to postulate a probability distribution of ε conditional on x and to integrate the decision rule over ε to obtain a form in terms of observable and structural parameters only. This turns out to be intractable, see Rust (1992).

The dimension of ε may vary with the number of elements in the agent's choice set, $D(x)$. We can identify each choice set as a set of integers: $D(x)=\{1,..., |D(x)|\}$, and let the decision space D be the set $D=\{1,..., \sup_{x \in X} |D(x)|\}$. Then whenever $|D(x)| < |D|$ we can consider the remaining components $|D| - |D(x)|$ as superfluous. Thus the vector needs to have at least as many components as the number of elements in $D(x)$ - see Rust (1994) for a full account.

Assumption 1 allows us to rewrite the value function 11 in the following way:

$$V(s) = \max_d [\tilde{v}(x, d) + \varepsilon(d)] \quad (14)$$

where

$$\tilde{v}(x, d) = \pi(x, d) + \beta \int V(s' | x, d) p(s' | x, d) \quad (15)$$

which represents the dynamic version of the static model discussed in Section 3.

Error structure: it may be useful at this point to summarize the error structure in the model. There are two sources of uncertainty which make the model stochastic:

1. A disturbance, which may be unpredictable or imperfectly predictable, is standard in models involving uncertainty about the future. Here it represents unexpected states occurring during the evolution of the system, e.g. exogenous shocks, both idiosyncratic and/or aggregate shocks. These are represented by the expected values of the variable regarding the future.
2. An error component which represents an "ignorance effect" arising in the case in which some state variables, also current, are imperfectly observable. We have assumed that the econometrician does not have perfect knowledge of all the state variables which each firm takes into

account when making current and future plans, in particular the information she has access to may have errors. Namely, the state variables which are perfectly known by the firm are random variables for the standpoint of the econometrician.

In what follows we deal with this second source of disturbance by reformulating the imperfect information case into a basic problem with perfect information. Intuitively, this may be done by defining a new system the state of which consists, at time t , of all variables the knowledge of which may be of benefit to the econometrician when modelling a decision observed in the data.

Our objective is to obtain the same structure as in Section 3 where we were able to concentrate on the difference between the pay-offs from two alternative choices. However, there is a fundamental difference between this model and the static structure discussed earlier. Since our problem is dynamic there is the possibility that the error terms are serially correlated. This is a likely case in many studies which use microdata where heterogeneity and persistence in idiosyncratic shocks cannot be easily neglected. If ε is autocorrelated, all future values of x are correlated with past and current values of ε . Bellman's equation 14 is carried out over a state space of expanding dimension and this makes estimation extremely difficult if not intractable. This problem is known in the literature as the "curse of dimensionality" (Bellman, 1957).¹⁵

In order to make the problem tractable we need to assume that the transition probability satisfies the conditionally independence assumption (Rust, 1987). Conditional Independence

We make the following assumption:

Assumption 2: *The transition probability for the process $\{x, \varepsilon\}$ may be expressed as:*

$$p(x', \varepsilon' | x, \varepsilon, d) = q(\varepsilon' | x') p(x' | x, d) \quad (16)$$

where $p\{x' | x, d\}$ is the transition probability of state x' given state x and decision d .

¹⁵Kwane and Wolpin (1994) propose a method for approximately solving discrete choice dynamic programming problems, based on simulation and interpolation. This method's partially solve the problem of Bellman's "curse of dimensionality" by obtaining approximate solutions for problems which would be otherwise intractable.

According to 16, conditional on the contemporaneous observed state variables, x 's, the unobservables do not depend on any previous state variables and decisions. Any serial dependence between ε and ε' is transmitted entirely through the observable state vector x' , which becomes a sufficient statistics for ε' . Moreover, conditional on the current state vector and decision, the probability density for the observable state variables x' , next period, depends on x and not on ε . Under conditional independence we may write equation 14 as:

$$V(s) = \max_d \left\{ \pi(x, d) + \varepsilon(d) + \beta \int [V(s' | x, d)] q(\varepsilon' | x') p(x' | x, d) \right\} \quad (17)$$

We have expressed the value function as an expectation which depends only on the observable state variables. In particular, consider the following conditional value function, ie the value function after having optimally chosen one of the available alternatives, for instance $d^* = 1$:

$$v_1(x) = \pi_1(x) + \varepsilon_1 + \beta \int \left[\max_j v_j(x') q(\varepsilon' | x') p(x' | x, d^* = 1) \right] \quad (18)$$

where $j=\{1, 2, \dots, J\}$ indicates the number of choices available to the firm and

$$v_j(x) \equiv E[V_j(s) | x, d^*(s) = 1] \quad (19)$$

is the value function in terms of only the observables. One of the very important consequences of Assumption 2, is that we may express the conditional value function in terms of the observable state variables.

It may be useful to summarize these last steps we have discussed. Starting from equation 17, we suppose that the firm's optimal choice at time t is $d=1$, for instance "invest" or "hire", which we observe in the data. Nevertheless, to model this decision we can only rely upon imperfect information and the conditional value functions, the value function relating to having optimally chosen "invest" or "hire" are, in principle, unknown functions of the state variables and structural parameters. However, we may rewrite the conditional value functions as expected value of functions of future paths of state variables and structural parameters, where the expectation is conditional on the current state variables and decisions, observable in the data, as indicated in equation 19 and implied by Assumption 2. It is essential to assume that, conditional on the observables, the autocorrelation in the

unobservables should be entirely captured by the autocorrelation in the observables.(see Aguirregabiria, 1997 for a discussion)

Using Assumptions 1 and 2 and some assumptions on the distribution of the unobservables, we may obtain a known functional form for equation 18. This will be play an important role in the estimate of the discrete choice model.

9 The Optimal Decision Rule

Suppose for simplicity that there are only two choices, $j=1,2$. The firm decides to choose option 1 if $v_1(x) > v_2(x)$ otherwise it chooses alternative 2. Using the definitions of the two conditional value functions we may express the optimal decision rule in terms of the differences in the expected pay-offs associated with the two alternative choices, as the following:

$$\delta(x, \varepsilon) = d^* = \begin{cases} 1 & \text{if } (\varepsilon_2 - \varepsilon_1) \leq (\pi_1(x) + v_1(x') - \pi_2(x) + v_2(x')) \\ 2 & \text{if } (\varepsilon_2 - \varepsilon_1) > (\pi_1(x) + v_1(x') - \pi_2(x) + v_2(x')) \end{cases} \quad (20)$$

$v_1(x)$ is specified as in equation 18 with an exact counterpart in the case of choice 2. The optimal discrete choice is, thus, expressed in terms of inequalities between the conditional value functions. In particular, equation 20 indicates that the firm decides to choose alternative 1 if the difference between the choice specific stochastic terms $(\varepsilon_2 - \varepsilon_1)$ is less than (or equal to) the difference between the conditional value functions relating to choice 1 and 2 respectively, and it will take action 2 if the reverse inequality holds. The same reasoning may be easily extended to the case of more than two choices. The decision rule takes the form of a threshold rule, where it is possible to compare the reward relative to each decision. In a static world, and hence not considering the terms $v_1(x')$ and $v_2(x')$ and assuming that the ε 's are distributed according a type I extreme value distribution we would have the static logit model discussed in Section 3.

In general, with J alternatives, given assumptions 1 and 2, the optimal discrete choice becomes:

$$d^* = \arg \max_{j \in \{1,2,\dots,J\}} \left(\pi^j(x) + \varepsilon^j(x) + E\beta v^j(x) \right) \quad (21)$$

The decision rule is expressed in terms of the observables. As we discussed

in the previous section, this is a consequence of Assumption 2 which implies that we can look for sufficient statistics which summarize all the information that is necessary for control purposes. In the absence of perfect knowledge on the state variables, the researcher may re-define the problem in terms of probabilistic states, making the problem conditional on the information available at current time, ie the observed decisions and the characteristics of the firms. There is a probability measure which allows us to redefine the basic problem with imperfect state information as a problem with perfect state information: the conditional choice probability. The problem may be stated in the following way: given the choice between 1 and 2 what is the probability that the we observe the employer selecting, say, choice 1, given the information available in the data?

10 Conditional Choice Probabilities

Assumption 2 allows us to adopt the same methodology as in the case of static discrete choice models. Define the following probability:

$$P(d | x) = \int I \{d = \delta(x, \varepsilon)\} q(d\varepsilon | x) \quad (22)$$

where $\delta(x, \varepsilon)$ is the decision rule, I is an indicator function and $q(\varepsilon | x)$ is the conditional distribution of ε given x , which we will define later. Probability $P(d | x)$ is the conditional choice probability ie the probability that the econometrician observes the firm optimally choosing $d = \delta(x, \varepsilon)$. It represents that sufficient statistics which summarizes all the information that is necessary for control purposes we were invoking in the previous sections.

Consider the conditional value function 18; using the definition in 22, we may rewrite the continuation value as a weighted average of the conditional value functions relating to each alternative choice where the weights are given by the probabilities that a decision is taken. Remember that the conditional value function is the value function conditional on having optimally chosen one of the available alternatives. Thinking again, for simplicity, of the two-choice case, we may rewrite the conditional value function relating to having optimally chosen alternative 1 as:

$$v_1(x) = \pi_1(x) + \varepsilon_1 + \beta \int \left[\sum_{j=1}^J P_j(x') v_j(x') p(x' | x, d^* = 1) \right] \quad (23)$$

Equation 23 is equivalent to formulation 18. Since the vector j comprises only two mutually exclusive choices: $j=1,2$, equation 23 becomes:

$$v_1(x) = \pi_1(x) + \varepsilon_1 + \beta \int [P_1(x') v_1(x') + P_2(x') v_2(x')] p(x' | x, d^* = 1) \quad (24)$$

In the case of two choices, the probability of making choice 1 may be written as (Hotz and Miller, 1993):

$$\begin{aligned} P_1 &= \int_{\varepsilon'_1=-\infty}^{\infty} \int_{\varepsilon'_2=-\infty}^{\varepsilon'_1+\Delta\bar{v}} dg(\varepsilon'_1, \varepsilon'_2) \\ &= \int_{-\infty}^{\infty} g_1(\varepsilon_1, [\varepsilon_1 + \bar{v}_2 - \bar{v}_1]) d\varepsilon_1 \end{aligned} \quad (25)$$

where $\bar{v}_j = \pi_j + v_j$, g is the error joint density. The integrand in the last line of 25 is the probability density for ε_1 , given decision rule 20, in particular $\bar{v}_1 + \varepsilon_1 \geq \bar{v}_2 + \varepsilon_2$ which applies when choice 1 is optimal. (McFadden, 1981).

In the general case, with J choices, the conditional choice probability satisfies the following:

$$\Pr\{d = \arg \max_{j \in \{1,2,\dots,J\}} (\pi^j(x) + \varepsilon^j(x) + E\beta v^j(x'))\} \quad (26)$$

In the next section we will show how choice probabilities allows us to represent the model in terms of the measured attributes of alternatives and individual characteristics, while maintaining the recursive structure inherent in the problem. Hence they may be directly used as representation of the intertemporal choice problem.

11 Reduction of the Model

The function v_d may be computed by solving the contraction fixed point which involve using computationally intensive backward recursion methods. Hotz and Miller (1993) provide an alternative representation which avoids having to solve explicitly for v_d . Their approach is based on a new representation of the valuation function which is expressed in terms of future pay-offs, choice probabilities and probability transitions of choices and outcomes that remain feasible in future periods.

Conditional on $d^*=j$, there is a relationship between the information error term, which captures the econometrician's imperfect knowledge on the state variables relevant for the firm's decision making and the conditional choice probability. In particular if $q(\varepsilon | x)$ follows a multivariate extreme value distribution then the conditional choice probability may be written as a multinomial logit

$$P(d | x) = \frac{\exp \{v_d(x)\}}{\exp \sum_i \{v_j(x)\}} \quad (27)$$

Consider the optimal decision rule in the case of two choices, 1 and 2, as in equation 20, where the optimal policy is defined in terms of inequalities between the conditional choice value functions. For each choice, $Q(\cdot)$ represents a positive, real-valued mapping from the differences in the conditional value functions.

$$\begin{aligned} P(d = 1) &= P\{(\varepsilon_2 - \varepsilon_1) \leq (\bar{v}_2 - \bar{v}_1)\} \\ &= Q(\bar{v}_2 - \bar{v}_1) = Q(\Delta \bar{v}) \end{aligned} \quad (28)$$

where $\Delta \bar{v}$ denotes the difference between the conditional valuation functions associated with the two alternatives choices, and $\bar{v}_1 = \pi_1 + v_1$. Hotz and Miller (1993, Proposition 1) show that $Q(\cdot)$ is a real valued, invertible function which allows us to express the difference in the conditional value functions in terms of the conditional choice probabilities as the following

$$\Delta \bar{v}(x) = Q^{-1}[P_1(x)] \quad (29)$$

In particular, if the information error term ε are i.i.d. extreme value distribution, we obtain the standard logit specification, according to which the difference in the conditional valuation functions, Δv , reduces to a log-odds transformation:

$$\Delta \bar{v}(x) = Q^{-1}[P_1(x)] = \log \left[\frac{1 - P_1(x)}{P_1(x)} \right] \quad (30)$$

We have, thus obtained the dynamic version of the structural multinomial logit model discussed in Section 3. Using this new expression of the valuation function we may now give a clearer representation to the firm's optimal decision process. Note that the term in square brackets in equation 24 may be rewritten as the following:

$$P_1(x') v_1(x') + (1 - P_1(x')) v_2(x') = v_1(x') + (1 - P_1(x')) (v_2(x') - v_1(x')) \quad (31)$$

Thus we equation 24 may be rewritten as

$$v_1(x) = \pi_1(x) + \varepsilon_1 + \beta \int \left\{ \log \left[\frac{1 - P_1(x)}{P_1(x)} \right] (1 - P_1) + P_1(x') v_1(x') \right\} p(x' | x, d^* = 1) \quad (32)$$

In the general case, consider the choice $j=1$ as the benchmark choice and express the problem in terms of difference in the conditional value functions, ie the log-odds ratios. We have that

$$\begin{aligned} v_d(x) &= \pi_d(x) + \beta \int \log \left[\sum_j \exp \{v_j(x') - v_1(x')\} \right] p(x' | x, d) \quad (33) \\ &\quad + \beta \int v_1(x') p(x' | x, d) \\ &= \pi_d(x) + \beta \int \log \left[\sum_j P_j(x') / P_1(x') \right] p(x' | x, d) \\ &\quad + \beta \int v_1(x') p(x' | x, d) \end{aligned}$$

The conditional value function comprises the one-period profit function, and the continuation value which is expressed in terms of the difference in the conditional choice probabilities relating to alternative choices.

The log-odds ratios may be estimated non-parametrically and constitutes an important part in the estimation of the discrete choice as we shall see later.

Rust (1994) notes that in the dynamic logit model does not suffer from the independence from irrelevant alternatives problem, which implies that the ratio of choosing any two alternatives does not depend on the attributes of the others. In the dynamic version, the way $v(\cdot)$ is specified implies that all alternatives are taken into accounts at each stage.

General Model

Let us summarize the results obtained in the previous sections by considering the general case in which there are J decisions. We may write the continuation term in the value function in general terms as:

$$\begin{aligned}
v^d(x') &= \int \sum_{j=1}^J P_j(x') v_j(x') q(\varepsilon'_j | x') p(x' | x, d^* = d) \\
&= \int \log \left[\sum_j P_j(x') / P_1(x') \right] p(x' | x, d) + \beta \int v_1(x') p(x' | x, d)
\end{aligned} \tag{34}$$

where $j=1$ is considered as benchmark choice. As discussed earlier, the expectation of the choice specific error terms, may be written in term of the vector of conditional choice probabilities, $P(x') \equiv \{P_1(x'), \dots, P_J(x')\}$. We only need a parametric specification of $v(\cdot)$ for a single alternative, $d=1$, instead for all alternatives, j . Given the benchmark conditional value function, the functional forms for the remaining conditional value functions are uncovered non-parametrically by inverting the mapping Q from value functions into choice probabilities. The invertible mapping reduces to a simple log-odds transformation when the error terms are i.i.d extreme value distribution because, in this case, function $q[\cdot]$ is a logit¹⁶. However, even if $P(\cdot)$ has a closed form in terms of the value functions $v_d(\cdot)$ if the ε 's are i.i.d. extreme, we need to know the functional form of the conditional choice value function $v_1(x')$. Usually we do not have a priori known functional form.

12 Closed Form Solution

The optimal decision rule may be written as

$$j = d^* = \delta^*(x, d) = \arg \max_{j \in \{1, 2, \dots, J\}} \left(\pi^j(x) + \varepsilon^j(x) + E\beta v^j(x') \right) \tag{35}$$

where v_j is the unique fixed point to the contraction mapping defined by the following condition¹⁷

¹⁶This distribution has the following form:

$q(\varepsilon | x) = \prod_d \exp \{ \{-\varepsilon + \gamma\} \exp \{-\exp \{-\varepsilon + \gamma\}\} \}$

The scale parameter γ , known as Euler's constant $\simeq 0.577$, shifts the extreme value distribution to obtain an unconditional mean equal to zero.

¹⁷The contraction mapping is $v: B \rightarrow B$, where B is the Banach space of all measurable bounded functions, $v: X \times D \rightarrow R$, where X is the state space and D is the decision space.

$$\begin{aligned}
v_j(x) &= \pi_j(x) + \varepsilon_j + \beta \int G(v_j(x') | x) p(x' | x, d) = \\
&\pi_j(x) + \varepsilon_j + \beta \int \int \max_j [v_j(x') | x] q(\varepsilon_j | x) p(x' | x, d)
\end{aligned} \tag{36}$$

and it is the unique solution to Bellman's equation 17. $G(\cdot)$ is the social surplus 6 defined in Section 4.

$$G(v_j(x') | x) \int \max_j [v_j(x') | x] q(\varepsilon_j | x) \tag{37}$$

For specific functional forms for the density $q(\varepsilon_j | x)$ we obtain very tractable forms of the conditional choice probabilities $P(d | x)$, the social surplus $G(v_j(x) | x)$ and the contraction mapping $v_j(x)$.

We follow Aguirregabiria (1999) and, given Assumptions 1 and 2, represent the optimal discrete choice by the following expression:

$$d^*(x; \theta) = \arg \max_j \{ \pi_j(x; \theta_\pi) + \varepsilon_j + v_j(\theta) \} \tag{38}$$

where $\theta = (\theta_\pi, \theta_f)$ and $v_j(\theta)$ has a closed-form solution. Exploiting the recursive structure of the problem, we may define the following Social Surplus:

$$G(x; \theta) \equiv E \left[\max_j \{ \pi_j(x; \theta_\pi) + \varepsilon_j + \beta v_j(x'; \theta) \} | x; \theta \right] \tag{39}$$

where θ is the vector of the parameters of the value function and θ_π is the vector of the parameters of the profit function we have omitted so far in order to simplify the notation. Specification 39 is equivalent to equation 6 in Section 12. Using expression 39 we may rewrite the continuation value as the following:

$$v_d(x'; \theta) = \int G(x'; \theta) p(x' | x, d; \theta_f) \tag{40}$$

and θ_f is vector of the parameters of the transition function and

$$G(x; \theta) = \sum_{j=1}^J P_j(x; \theta) \left[\pi_j(x; \theta_\pi) + q_d(P[\theta]) + \beta \int G(x'; \theta) p(x' | x, j; \theta_f) \right] \tag{41}$$

Equation 41 defines $G(x; \theta)$ as the fixed point of a contraction mapping. Numerical solutions of this contraction mapping may be obtained for particular

values of θ . Methods such as the nested solution-estimation algorithm have successfully estimated discrete choice structural models, see Rust (1994) and the survey by Eckstein and Wolpin (1989). These methods require an explicit solution of the model at each iteration in the search for the parameter estimates. This may become very demanding from a computational point of view. For instance, Rust (1994) suggests a nested fixed point algorithm: an "inner" contraction fixed point algorithm to compute $v(\theta)$ for each trial value of θ , and an "outer" hill-climbing algorithm to search for the value of θ that maximizes full information maximum likelihood, $L(\theta)$.

However, we can avoid having to solve the fixed point contraction mapping since $G(\cdot)$ has a closed form solution which may be expressed in terms of the one-period profit function, the conditional choice probabilities and the transition probabilities. To obtain this we need to discretize of the space of the observed state variables. Let $\{x^1, x^2, \dots, x^M\}$ be a discretization grid. Using this discretization we may write the contraction mapping in matrix form as the following:

$$G(\theta) = \sum_d P_d(\theta) * [\pi_d(\theta_\pi) + q_d(P[\theta]) + \beta F_d(\theta_f) G(\theta)] \quad (42)$$

where $v(\theta)$, $P_d(\theta)$ and $q_d(P[\theta])$ are $M \times 1$ vectors, $*$ is the element by element product; $\pi_d(\theta_\pi)$ is an $M \times (J-1)$ matrix and $F_d(\theta_f)$ is the $M \times M$ matrix of transition probabilities conditional on discrete choice d . The solution to the contraction mapping 42 is given by

$$G(\theta) = (I_M - \beta F(\theta))^{-1} \left\{ \sum_{j=1}^J P_d(\theta) * (\pi_d(\theta_\pi) + q_d(P[\theta])) \right\} \quad (43)$$

where $F(\theta) = \sum_{j=1}^J P_d(\theta) F_d(\theta_f)$ is the matrix of unconditional transition probabilities; $P_d(\theta)$ is a diagonal matrix with the conditional choice probabilities on the diagonal. Let θ^* be the true vector of structural parameters in the population, and define $P_d^* \equiv P_d(\theta^*)$, $F_d^* \equiv F_d(\theta^*)$, $F^* \equiv F(\theta^*)$, where the asterisk indicates the true parameter. Therefore, for any θ close enough to θ^* we have that

$$G(\theta) \approx (I_M - \beta F^*)^{-1} \left\{ \sum_{j=1}^J P_d^* * (\pi_d(\theta_\pi) + q_d(P^*)) \right\} \quad (44)$$

Then 41 may be written in the following estimable form:

$$v_d(\theta) = F_d^* (I_M - \beta F^*)^{-1} \left\{ \sum_{j=1}^J P_d^* * \pi_d(\theta_\pi); \sum_{j=1}^J P_d^* * q_d(P^*) \right\} \quad (45)$$

Equation 45 is a closed-form solution to problem 34 and we may obtain consistent estimates of the structural parameters without having to solve directly the structural model. Indeed, we know the form of the one-period profit functions $\pi(\theta_\pi)$ and we can estimate F^* , F^{d*} and P^* non-parametrically.

13 Estimation

Semi-parametric estimation of model 45 proceeds in 2 stages:

Stage 1.

- Nonparametric estimation of the conditional choice probabilities for each discrete alternative over the range of non-discretized state variables. The local probability P_j of making choice j conditional on the vector ξ of state variables is estimated as

$$P_j = \frac{\frac{1}{nh} \sum_{i=1}^n K\left(\frac{\xi - x_i}{h}\right) d_{ij}}{\frac{1}{nh} \sum_{i=1}^n K\left(\frac{\xi - x_i}{h}\right)}$$

where n is the total number of observations, $K(\cdot)$ is the kernel (local averaging) function, typically chosen as Gaussian, and h is the window size, typically taken as proportional to n^{-6} .

- Estimation of the transition probabilities of the observed state variables to obtain estimates for θ_f . For this purpose, we require a discretization of the observed state variables. The optimal size of the discretizing grid depends on the number of state variables and the dimension of the dataset. Grid size should be chosen to ensure approximately equal numbers of observations in each cell. When the state variable dimension exceeds one, the need to avoid sparsely populated cells can force a coarse mesh. Transition probabilities are estimated as the sample frequencies at the mean cell values of the state variables. In general the matrix of s -step-ahead probabilities is $p(x_{t+s} | x_t, d = j) = F^{s-1} F_j$, $s > 0$, where F

is the $M \times M$ matrix of transition probabilities and F_j is the $M \times M$ matrix of one-period transition probabilities conditional on discrete choice $d=j$. In a stationary Markov environment, all future probabilities can be obtained from the one-step-ahead probabilities, see Aguirregabiria (1993).

Stage 2.

- Having estimated the conditional choice probabilities and the transition probabilities in Stage 1, for any values of the parameters in θ_π , we can construct the values for expression 45

$$v_d(x, \theta) = \mathbf{F}_d^* (\mathbf{I}_M - \beta \mathbf{F}^*)^{-1} \left\{ \sum_{j=1}^J P_d^* * \Pi_d(\theta_\pi); \sum_{j=1}^J P_d^* * q_d(P^*) \right\} \quad (46)$$

where $\Pi_d(\theta_\pi)$ is the matrix of the one-period profit functions and P_d^* is the matrix of the conditional choice probabilities. Under the hypothesis that the error terms ε 's are i.i.d. extreme value, this enables us to calculate the adjustment probabilities

$$P_{(d)}(v; \theta_\pi) = \frac{\exp \{v_d(x; \theta_\pi)\}}{\sum_j \exp \{v_j(x; \theta_\pi)\}} \quad (47)$$

Given a set of orthogonal variables Z , satisfying

$$E \left(Z \left[I(d=j) - P_{(d)}(\Pi, \hat{v}; \theta_\pi) \right] \right) = 0$$

we may define the vector e of generalized residuals by $e(\theta_\pi) = I(d=j) - P_{(d)}(\Pi, \hat{v}; \theta_\pi)$. This motivates the generalized method of moments (GMM) estimator obtained by minimization of the criterion function

$$\min_{\theta} e' Z (Z' \Omega Z)^{-1} Z' e \quad (48)$$

where Z is the matrix of instrumental variables, $\Omega = E(ee')$. Because Ω is unknown and unknowable, we follow White's (1980) two stage procedure. At the first stage Ω is replaced by the identity matrix and

at the second stage $W = (w_{ij}) = Z'\Omega Z$ is calculated from the first stage residuals \hat{e} by

$$w_{ij} = \sum_{k=1}^n z_{ki} z_{kj} \hat{e}_k^2 \quad (49)$$

Hotz and Miller (1993) show that the conditional choice probability estimator is consistent and asymptotically normal.

14 An Example

To illustrate this, consider again the simplified case in which the firm has to decide between two mutually exclusive choices. We follow Rota (2000) and analyze firms' intertemporal employment decisions in the presence of lumpy adjustment costs.¹⁸ In a discrete time, infinite horizon framework, the firm's intertemporal problem is to find an optimal employment policy. In the presence of lumpy and kinked adjustment costs, firms may find it optimal not to vary the number of employees, even in the presence of shocks. Firms only make those changes in the labour input which are justified by sufficiently large departures of desired employment from their most recent choice of the number of employees. The adjustment process is lumpy and intermittent: in the face of a shock, a firm may decide that it is optimal to maintain the same number of employees and to postpone adjustment to the future; there is the possibility of corner solutions in the demand for labour. In this case, the standard marginal conditions for optimality given by the Euler equation fail to hold. At the beginning of period t , the firm chooses the profit maximizing level of employment. It observes the beginning of period number of workers L_{t-1} , the current wage, W_t and the productivity shock, ω_t , but it is uncertain about future wages and productivity shocks. There are two exogenous state variables in the model: wages and productivity shocks. Wages are imperfectly observed by the econometrician due to lack of information on the costs of

¹⁸In Del Boca and Rota (1998) the analysis of the legal provisions on hiring and firing in Italy show how that the legislation favours the use of collective firings as a way of reorganizing personnel, while the strict unfair dismissal regulation makes discharges very difficult. The implied pattern is alternating regimes of large adjustment and of non-adjustment, consistent with that generated by fixed adjustment costs. See Hamermesh (1989) for the seminal work on labour demand in the presence of fixed costs.

Moreover, the absence of a system of severance payments in Italy, makes alternative cost structures, such as linear costs less appropriate.

monitoring workers, and heterogeneity with respect to human capital. Thus the state vector, s , comprises an unobservable component, ε , which captures the heterogeneity in both the outcome paths and the employment choices observed in the data:

$$s \equiv [L_{t-1}, W_t, \omega_t, \varepsilon_t] \quad (50)$$

The vector of observables is:

$$x \equiv [L_{t-1}, W_t, \omega_t] \quad (51)$$

The one-period profit function has a very simple form:

$$\Pi_t = \{AL_t^\alpha - b \mid \Delta L_t \mid -K1(L_t \neq L_{t-1}) - W_t L_t\} \quad (52)$$

where $A > 0$, $0 < \alpha < 1$, $b > 0$ is the cost per unit of labour, K is the fixed cost and $1(L_t \neq L_{t-1})$ is an indicator function. When the firm adjusts employment it pays both kinked and lump-sum costs. In this model adjustment costs are symmetric.

Under assumptions 1 and 2, the conditional value function, relating to choice $d=A$, ie adjust the number of employees has the following form:

$$\begin{aligned} v_A(x) = & \pi_A(x) + \varepsilon_A + \\ & \beta \int \{v_A(x') + [1 - P_A(x')] [v_{NA}(x') - v_A(x')]\} p(x' \mid x, d^* = A) \end{aligned} \quad (53)$$

where the conditional value function is expressed in terms of the observable state variables and, since there are only two choices, it must be that $P_A = (1 - P_{NA})$. Equation 53 is the analogous of equations 23 and 24 in Section 11. We may write the value function conditional on having optimally chosen alternative $d=NA$, as

$$\begin{aligned} v_{NA}(x) = & \pi_{NA}(x) + \varepsilon_{NA} + \\ & \beta \int \{v_A(x') + [1 - P_A(x')] [v_{NA}(x') - v_A(x')]\} p(x' \mid x, d^* = NA) \end{aligned} \quad (54)$$

Equation 53 may be rewritten as:

$$v_A(x) = \pi_A(x) + \varepsilon_A + \beta \int \left\{ v_A(x') + [1 - P_A(x')] \gamma \log \left[\frac{1 - P_A(x')}{P_A(x')} \right] \right\} p(x' | x, d^* = A) \quad (55)$$

where γ is a constant of proportionality. The conditional value function for decision $d^* = NA$ is defined symmetrically.¹⁹

Before proceeding to the estimation of the discrete choice implied by the existence of non-convex adjustment costs, it is worth noticing the following property of our model. If we concentrate on the value function conditional on adjusting, ie conditional on the firm deciding to change from L_{t-1} to L_t , we may take the first order condition with respect to L_t and obtain the following:

$$\frac{\partial v_t^A}{\partial L_t} = \frac{\partial \pi_t^A}{\partial L_t} + \beta \gamma_{t+1} E \left[\frac{\partial P_{t+1}^{NA}}{\partial L_t} q_{t+1} + P_{t+1}^{NA} \frac{\partial q_{t+1}}{\partial L_t} \right] = 0 \quad (56)$$

where $q_{t+1} \equiv \log \left[\frac{1 - P_{t+1}^A}{P_{t+1}^A} \right]$. Here, the first term is standard and relates to increased profits from the current period adjustment, but this is augmented by the square-bracketed term which relates to the consequences of maintaining this new level of employment in subsequent periods, as the result of lumpy adjustment costs. Using the one-period profit function 52, we obtain the following marginal productivity condition:

$$\left(\frac{Q}{L} \right)_t = \Psi_1 W_t + \Psi_2 (Pq)_{t+1} + v_{t+1} \quad (57)$$

where

$$(Pq)_{t+1} = \left[\frac{\partial P_{t+1}^{NA}}{\partial L_t} q_{t+1} + P_{t+1}^{NA} \frac{\partial q_{t+1}}{\partial L_t} \right] \quad (58)$$

and v_{t+1} is a realization error. The structural parameters are related to the coefficients of 57 by:

¹⁹The decision rule is the following:

$$\delta(x, \varepsilon) = d^* = \begin{cases} 1 & \text{if } (\varepsilon_{NA} - \varepsilon_A) \leq (\pi_A(x) + v_A(x') - \pi_{NA}(x) + v_{NA}(x')) \\ 2 & \text{if } (\varepsilon_{NA} - \varepsilon_A) > (\pi_A(x) + v_A(x') - \pi_{NA}(x) + v_{NA}(x')) \end{cases}$$

$$\Psi_1 = \frac{1}{\alpha}; \Psi_2 = \frac{b}{\alpha}; \Psi_3 = -\frac{\beta\gamma}{\alpha}; \quad (59)$$

Equation 57 is a simple marginal productivity relation augmented by the forward-looking term $(Pq)_{t+1}$ which captures the alternatives of future adjustment or non- adjustment arising from the presence of fixed adjustment costs.²⁰

Equation 57 may be estimated using a two stage procedure; we first estimate the conditional choice probabilities P_{t+1}^A and P_{t+1}^{NA} and the slopes $\frac{\partial P_{t+1}^{NA}}{\partial L_t}$ non-parametrically, using a kernel method, and construct the term $(Pq)_{t+1}$. In the second stage we use GMM to estimate the augmented marginal productivity equation 57, by conditioning the model on adjustment occurring at time t. We thus consider the sub-sample of interior solutions, which comprises the firms which have optimally adjusted at time t.²¹

The semi-reduced form 57, obtained by exploiting the first order conditions for optimality, allows us to obtain estimates of the parameters of the production function. The important new result is that, conditional on having adjusted employment between t-1 and t, the marginal revenue product of labour will differ from the wage the firm pays by a forward looking term which embodies the discrete choices implied by the cost structure. This procedure allows us to exploit standard orthogonality conditions in estimation and it may also provide a way of properly estimating intertemporal relations for a variety of other problems (eg investment demand and inventory).

However, the fixed cost parameter K is not identified from these estimates. In order to estimate fixed costs we need to directly model the discrete choice. In order to do this, we need to refer to the methodology proposed by Hotz and

²⁰Suppose that the firm increases (decreases) employment in period t but, because of high adjustment costs, expects not to adjust in the subsequent period ($P_{t+1}^{NA} > P_{t+1}^A$). Thus P_{t+1}^{NA} is high and q_{t+1} positive. Suppose that at the beginning of period t+1, L_t is above the desired level, then the derivatives of both P_{t+1}^{NA} and q_{t+1} with respect to L_t will typically be positive. This implies that the term q_{t+1} will be positive. The firm expects to reduce employment in the future but, at present, L_t is higher than it would have obtained in the absence of adjustment costs, and marginal productivity is low relative to the wage rate. In the opposite case, the firm intends to hire new employees in the future, but at present L_t is below the desired level and marginal productivity is high with respect to the wage rate (the term q_{t+1} is negative). The term q_{t+1} , therefore allows us to take into account the future choices the company faces subsequent to adjustment at t.

²¹For this purpose, we adopt the following selection window: we keep or drop the observation on L_t according to whether it is different from or equal to L_{t-1} .

Miller (1993) and Aguirregabiria (1999) and discussed in Section 15. Nevertheless, we are able to simplify their procedure by utilizing the structural parameter estimate obtained from estimation of the first order condition 57.

To estimate the size of fixed costs we need to estimate first the transition probabilities p_A and p_{NA} . To estimate the discrete decision process we require a discretization of the space of the observable exogenous state variables, W_t and ω_t . We consider a grid of $m=52=25$ cells with approximately equal numbers of observations. Define:

$$\begin{pmatrix} v_A(\tilde{x}) \\ v_{NA}(\tilde{x}) \end{pmatrix} = (\mathbf{I}_{2m} - \beta \mathbf{R})^{-1} \begin{pmatrix} \pi_A(\tilde{x}) + \varepsilon_A \\ \pi_{NA}(\tilde{x}) + \varepsilon_{NA} \end{pmatrix} \quad (60)$$

where

$$\mathbf{R} = \begin{bmatrix} \mathbf{T}_A \mathbf{P}_A & \mathbf{T}_A \mathbf{P}_{NA} \\ \mathbf{T}_{NA} \mathbf{P}_A & \mathbf{T}_{NA} \mathbf{P}_{NA} \end{bmatrix} \quad (61)$$

and where \mathbf{I} is the identity matrix, \mathbf{T}_A and \mathbf{T}_{NA} are $(m \times m)$ matrices of the one-period transition probabilities, \mathbf{P}_A and \mathbf{P}_{NA} are diagonal matrices with the conditional choice probabilities on the diagonal, and \tilde{x} indicates the discretized state space. The value functions and profits are evaluated at the mean cell values of the state variable. The one-period transition probabilities are estimated non-parametrically.

The vector of generalized residuals is

$$e_{it} = d_{it} - \Phi\left(\frac{J_{it}}{\sigma_J}\right) - \delta \quad (62)$$

where J_{it} indicator function, Φ is the standard normal distribution function, σ_J is the sample standard deviation of the J_{it} and δ is an intercept. According to the structural methodology, firm i should adjust in period t if $J_{it} > 0$ but otherwise it should remain with an unchanged labour force.

In the first stage we choose θ to minimize

$$e'Z(Z'\Omega Z)^{-1}Z'e \quad (63)$$

where Z is a matrix of instruments, Ω is a covariance-weighting matrix which corrects for heteroscedasticity (White, 1980). At the first stage Ω is replaced by the identity matrix and at the second stage $Z'\Omega Z$ is calculated using the first stage residuals.

It should be noticed that fixed costs parameters are not only associated with the constant terms in the choice specific intertemporal profits. Fixed costs are also associated with some components of v_{it}^d . It is crucial for the identification of these parameters that these components present significant sample variation. Estimation on a sample of more than 3000 Italian firms recorded over the period 1982-1989 indicates that adjustment costs are very high. The median level of fixed costs corresponds to approximately 3.65 times average unit labour costs.

15 What Do We Gain By Being Discrete?

In this paper we have emphasized how the DDP model represents a natural framework within which to analyze agents' behaviour when optimal inaction generates censoring in observed decisions. Certain phenomena are naturally discrete - we have mentioned Rust's machine replacement model or Hotz and Miller's fertility decisions. It is obviously desirable to model these as discrete. Considering individual agents decisions overcomes the representative agent's problem, as the McFadden's example explains: half the population travels always by bus and half always by car, while the "representative person" travels 50% by bus and 50% by car. Is the latter truly representative?

Other economic problems may result in agents behaving in a discontinuous manner, alternating period of actions to ones in which their optimal decision is "doing nothing". One of the many examples, is given by the optimizing response of a firm in the presence of fixed adjustment costs; the firms may optimally choose to postpone the decision of changing one or more inputs to the future even in the presence of shocks. This behaviour generates corner solutions in the optimal adjustment path and a break in the standard first order conditions given by the Euler equation. One solution would be to consider only the subset of interior solutions and recover the Euler equations. However, we would lose the information contained in the discrete decision, and obtain inefficient parameter estimates. In order to exploit the information contained in the discrete choice we need to model this behaviour as a discrete (or mixed continuous-discrete) decision process. Longitudinal datasets, which record cross sections over time, show that often agents do not respond to observed changes in the variables relevant for their decision making, even when these are of a large magnitude.

It is tempting to suppose that aggregation makes discrete decisions con-

tinuous, but many studies have found that microeconomic frictions may influence aggregate dynamics. The effect of heterogeneity of agents, idiosyncratic uncertainty and lack of coordination can imply highly nonlinear behavior in aggregate time series. Hence, reference to microeconomic models becomes relevant if aggregate dynamics are to be correctly understood. Caballero and Engel (1992, 1993 and 1994) propose an adjustment hazard function which determines the probability that a unit adjusts, in a given time interval, as a function of the difference between the current and the target value of a variable (inventories, durables, capital etc.). The model encompasses both (S,s) rules (where the hazard functions jumps from zero to infinity) and the linear-quadratic model (constant hazard). Only in the second case does the model generate aggregate dynamics which are indistinguishable from those of the quadratic cost representative agent model, while a non constant hazard introduces nonlinearities and complex dynamics into aggregate relationships. In this case, the propensity to change the stock of inputs is positively related to the size of the shock. The response to an aggregate shock will depend on whether most units are within regions with large or small adjustment hazards.

There is a fundamental methodological difference between discrete choice and continuous choice models. In the former class of models, choice can be varied continuously and the net effect of a marginal change on costs and benefits helps to characterize the optimal decision: for example, concavity of the objective function, implies that the solution to a first-order condition fully characterizes the optimal choice. In the latter class of models, when we deal with a finite number of choices, the value of taking each choice must be compared with the others, so optimal behaviour is characterized in terms of inequalities. Modelling the choice set as finite implies that people with similar but not identical situations might make the same choice. See, for instance McFadden (1981) in the case of the Social Surplus Function and Manki's (1993) in his static path representation according to which a basic social feature of decision making is that one learns from the experience of others and emulates them. Information about their marginal differences is lost and by only focusing on broad based options, the econometrician must consequently maintain stronger assumptions about the functional forms of preferences, opportunities and the nature of unobserved heterogeneity throughout the sample population in order to achieve econometric identification. However, it is also possible to apply the DDP framework to models where the choice variables come from a continuous support; indeed one method of handling a

continuous choice model is to first discretize it with a finite number of states and choices.

One final reason should motivate the study of certain economic phenomena as DDP. While up until quite recently the estimation of dynamic structural models has been computationally very demanding (Flinn and Heckman (1982) avoid this by relying upon a highly specialized structure), the method discussed in this paper, proposed by Hotz and Miller, allows us to obtain a closed form solutions in terms of conditional choice probability, transition probabilities and the one-period reward function. This greatly reduces the computational burden implied by ML estimation techniques, which until quite recently it was believed to be the only way to estimate dynamic structural discrete choice problems.

All the structural estimation methods we have discussed, including that proposed by Hotz and Miller, still rely upon the simplifying assumption of conditional independence: this is needed in order to obtain a tractable econometric model. If the model has serially correlated choice specific error terms, lagged values of these error components would enter as additional state variables in the value functions. This may generate an increasingly expanding state space and significantly aggravate the computational burden of solving the DDP model. While it is plausible that the error terms are serially correlated, there is no method available which allows us to relax the conditional independence restriction.

As a final encouragement to the use of discrete models of behaviour I return to one of the earliest but also most literary contributions to this literature: "Be not too tame neither, but let your own discretion be your tutor: suit the action to the word, the word to the action" (W. Shakespeare, Hamlet).

Appendix 1 McFadden's Probabilistic Choice Model

Probabilistic choice system describes the observable distribution of demands by a population of consumers. The distribution of demands in a population is the result of individual preference maximization, with preferences influenced by unobserved variables. What are the features of the observable distribution of demands that are necessary or sufficient for their consistency with the hypothesis of random preference maximization. Stochastic version of the theory of revealed preference.

A choice probability specifies the probability of choosing an action $i \in \mathbf{I}$, where \mathbf{I} is a set indexing alternatives, given that a selection must be made from the finite choice set $\mathbf{B} \in \mathcal{B}$ and that the decision-maker has characteristics $\mathbf{s} \in \mathbf{S}$. Choice probabilities are nonnegative and sum to one, they depend only on the measured attributes of alternatives and individual characteristics. Consider a probability measure μ depending on $\mathbf{s} \in \mathbf{S}$ on the space of utility functions on \mathbf{I} . Then μ gives the distribution of tastes in the population of individuals with characteristics $\mathbf{s} \in \mathbf{S}$. Then

$$\begin{aligned} P(i_1 | \mathbf{B}, \mathbf{s}) &= \int_{u_1=-\infty}^{\infty} \int_{u_2=-\infty}^{u_1} \dots \int_{u_n=-\infty}^{u_1} f^B(u_1, \dots, u_n; \mathbf{s}) du_1 \dots du_n \quad (64) \\ &= \int_{u=-\infty}^{\infty} F_1^B(u, \dots, u; \mathbf{s}) du \end{aligned}$$

where F^B is the cumulative distribution function of f^B , and F_1^B denotes the derivative of F^B with respect to its first argument. Choice is determined by utility maximization and there is almost always a unique utility-maximizing alternative. Is the model of individual utility maximization identifiable from the observed distributions of demands. Could other model generate the same observations? A necessary and sufficient condition for consistency with random preference maximization, analogous to the strong axiom of revealed preferences has been established by McFadden and Richter (1970). Strong axiom of stochastic revealed preference states that for any finite sequence of trials $(\mathbf{B}^1, \mathbf{C}^1), \dots, (\mathbf{B}^M, \mathbf{C}^M)$, with repetition permitted,

$$\sum_{m=1}^M P(\mathbf{C}^m | \mathbf{B}^m, \mathbf{s}) \leq N((\mathbf{B}^1, \mathbf{C}^1), \dots, (\mathbf{B}^M, \mathbf{C}^M)) \quad (65)$$

where $N((\mathbf{B}^1, \mathbf{C}^1), \dots, (\mathbf{B}^M, \mathbf{C}^M))$ is the maximum number of successful trials in the sequence consistent with a single preference order. This condition is necessary but not sufficient.

Appendix 2 Discrete and Continuous Choices

We may partition choices into two types: continuous and discrete choices. Define continuous choices made on date t by the M -dimensional vector $c_t = (c_{1t}, \dots, c_{Mt}) \in C$. Decisions made at t only exploit information available at that time. We may write $c_t = c(s_t)$ or $c_{mt} = c(s_{mt})$, where s represents the state of the world at time t . Similarly, define discrete choices as $d_{kt} \in \{0, 1\}$. The firm pick one the several actions $k \in \{0, \dots, K\}$ in each period t . If a choice is selected then $d_{kt} = 1$. Information plays the same role as in the case of continuous choices, hence $d_{kt} = d(s_{kt})$. We assume that choices are mutually exclusive so that $d_t \equiv (d_{1t}, \dots, d_{Kt})$ fully summarize the firm's discrete choices:

$$\sum_{k=0}^K d_{kt} = 1 \quad (66)$$

The firm's preferences are defined over the choices it makes over her lifetime, and over the state variables. Thus s evolves according a stochastic law of motion which represents the subjective beliefs about future state variables. In the case in which s is discrete, ie it has finite support, we have the following probability distribution:

$$\Pr\{s_{t+1} \mid s_t, c_t, d_{kt} = 1\} \equiv F_{kt}(s_{t+1} \mid s_t, c_t) \quad (67)$$

The future state variable s_{t+1} is randomly determined by the current state variable s_t and current choices (c_t, d_t) . This equation is also assumed to underlie the data generating process.

The stream of profits may be thus defined as

$$E \left[\sum_{t=0}^T \sum_{k=0}^K d_{kt} \pi_{kt}(c_t, s_t) \mid z_0 \right] \quad (68)$$

where π_{kt} is a real valued mapping from CxS. Whatever the date/state coordinate pair is (t, s) the firm makes continuous choice $c(s_t)$ and discrete choice $d(s_t)$. In particular, if the firm's state at time t is s_t , it might choose j , ie set $d_{jt} = 1$, and make continuous choice c_t . The the profit it would receive that period is $u_{jt}(c_t, s_t)$. Integrating over $s \in S$, conditional on her initial state s_0 , her expected profit at time t is

$$E \left\{ \sum_{k=0}^K d_{kt} \pi_{kt}[c_t, s_t] \mid s_0 \right\} \quad (69)$$

In an infinite time horizon we have:

$$E[\sum_{t=0}^{\infty} \sum_{k=0}^K d_{kt} \beta^t \pi_{kt}(c_t, s_t) \mid s_0] \quad (70)$$

firm's problem: for each $t \in \{0, 1, \dots, T\}$ and $s \in S$, it picks $(c_t, d_t) \in C \times D$ to maximize 70 subject to 66. The solution to this problem can be interpreted as an optimal decision rule $(c_t(s_t), d_t(s_t))$. We may express the problem in terms of the following value function:

$$v(s) = \sum_{k=0}^K d_k [\pi_k(s, c) + \beta \int v(s') dF_k(s' \mid c, s)] \quad (71)$$

where, because of the stationarity. of the model, the time subscript has been dropped and the apostrophe indicates variables not known at the present time. All the dynamic factors are transmitted through the transition probability $F_k(s' \mid c, s)$. This problem can be further decomposed: for each $k \in \{0, \dots, K\}$, the firm chooses c_t to maximize

$$\pi_k(c, s) + \beta \int v(s') dF_k(s' \mid c, s) \quad (72)$$

Let c^k denote the solution to maximizing 72, and successively substitute the solution to the K subproblems into current profit and transition probabilities. The define the reduced +form profits $\pi_k \equiv \pi_k(c^k, s)$ and transition probabilities $F_k(s' \mid s) \equiv F_k(s' \mid c, s)$. The optimal discrete choice d^* maximizes

$$\pi_k(s) + \int v(s') dF_k(s' \mid s) \quad (73)$$

over $k \in \{0, \dots, K\}$. We have thus obtained a discrete choice problem ²².

²²Pakes considers both continuous and discrete controls.

References

- ABEL, A. and EBERLY, J. (1996), "Investment and q with Fixed Costs: An Empirical Analysis", *manuscript*.
- AGUIRREGABIRIA, V. (1997), "Estimation of Dynamic Programming Models with Censored Dependent Variables", *Investigationes Economicas*, 21, 167-208.
- AGUIRREGABIRIA, V. (1999), "The Dynamics of Markups and Inventories in Retailing Firms", *Review of Economic Studies*, 66, 275-308.
- AGUIRREGABIRIA, V. and ALONSO-BORREGO, C. (1999), "Labor Contracts and Flexibility: Evidence from a Labor Market Reform in Spain", *manuscript*.
- ATTANASIO, O. (1997), "Consumer Durables and Inertial Behavior: Estimation and Aggregation of (S,s) Rules for Automobile Purchases", *manuscript*.
- BAR-ILAN and BLINDER (1992), "Consumer Durables and the Optimality of Usually Doing Nothing", *Journal of Money, Credit and Banking*, 24, 253-272.
- BARNETT, S. and SAKELLARIS, P. (1998), "Nonlinear response of firm investment to Q : testing a model of convex and non-convex adjustment costs", *Journal of Monetary Economics*, 42, 261-288.
- BENTOLILA, S. and BERTOLA, G. (1990), "Firing Costs and Labor Demand: How Bad is Eurosclerosis?", *Review of Economic Studies*, 57, 381-402.
- BERTOLA G. and CABALLERO, R. (1994) "Irreversibility and Aggregate Investment", *Review of Economic Studies*, 61, 223-246.
- BERTSEKAS, D. (1987), "Dynamic Programming. Deterministic and Stochastic Models", (Englewood Cliffs, N.J.: Prentice-Hall)
- BLACKWELL, D. (1962), "Discrete Dynamic Programming", *Annals of Mathematical Statistics*, 33, 719-726.
- BLACKWELL, D. (1965), "Discounted Dynamic Programming", *Annals of Mathematical Statistics*, 36, 226-235.
- CABALLERO, R. (1997), "Aggregate Investment", *NBER Working Paper* no. 6264.
- CABALLERO, R. and ENGEL, E. (1992), "Beyond Partial Adjustment", *American Economic Review*, Papers and Proceedings, 82, 360-364.
- CABALLERO, R. and ENGEL, E. (1993), "Microeconomic Adjustment Hazards and Aggregate Dynamics", *Quarterly Journal of Economics*, 108, 359-383.

- CABALLERO, R. and ENGEL, E. (1994), "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach", NBER Working Paper no. 4887.
- COSTANTINIDES, G. and RICHARD, S. (1978), "Existence of Optimal Simple Policies for Discounted-Cost Inventory and Cash Management in Continuous Time", *Operations Research*, 26:620-636.
- COOPER, R., HALTIWANGER, J. and POWER, L. (1995), "Machine Replacement and the Business Cycle: Lumps and Bumps", *NBER Working Paper* no. 5260.
- CRAMER, J. (1991), *"The Logit Model"*, (New York: Chapman and Hall).
- DEL BOCA, A. and ROTA P. (1998), "How Much Does Hiring and Firing Cost? A survey study on a sample of Italian firms," *Labour*, 12, 427-449.
- DEBREU, G. (1960), "Review of R. D. Luce Individual Choice Behavior", *American Economic Review*, 50, 186-188.
- DIXIT, A. and PINDYCK, R. (1994), *Investment under Uncertainty*, (Princeton, New Jersey: Princeton University Press).
- EBERLY, J. (1994), "Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases", *Journal of Political Economy*, 31, 403-436.
- ECKSTEIN, Z. and WOLPIN, K. (1989), "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models", *Journal of Human Resources*, 24, 562- 598.
- FAFCHAMPS, M. and PENDER, J. (1997), "Precautionary Saving, Credit Constraints, and Irreversible Investment: Theory and Evidence from Semi-arid India", *Journal of Business & Economic Statistics*, 15, 180-194.
- FLINN, C. and HECKMAN, J. (1982), "New Methods for Analysing Structural Models of Labor Force Dynamics", *Journal of Econometrics*, 18, 115-168.
- GROSSMAN, S. and LAROQUE, G. (1990), "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods", *Econometrica*, 58,25-51.
- HOPENHAYN, H. and ROGERSON, R. (1993), "Job Turnover and Policy Evaluation: A General Equilibrium Analysis", *Journal of Political Economy*, 101, 915-938.
- HAMERMESH, D. (1989), "Labor Demand and the Structure of Adjustment Costs," *American Economic Review*, 79, 675-689.
- HAMERMESH, D. and PFANN, G. (1994), "Adjustment Costs in Factor Demand," *Journal of Economic Literature*, 34, 1264-1282.

- HOTZ, J. and MILLER, R. (1993), "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497-529.
- HOTZ, J., MILLER, R., SANDERS, S. and SMITH, J. (1994), "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, 61, 265-289.
- KEANE M. and K. WOLPIN (1994), "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence", *Review of Economic and Statistics*, 76, 648-672.
- KEANE M. and K. WOLPIN (1997), "Introduction to the JBES Special Issue on Structural Estimation in Applied Microeconomics", *Journal of Business & Economic Statistics*, 15, 111-114.
- KENNET, M. (1993), "Did Deregulation Affect Aircraft Maintenance? An Empirical Policy Analysis", *Rand Journal*, 24, 542-558.
- LAM, P. (1991), "Permanent Income, Liquidity, and Adjustments of Automobile Stocks: Evidence from Panel Data", *Quarterly Journal of Economics*, 106, 203-30.
- LUCE, R. (1959), *Individual Choice Behavior: A Theoretical Analysis*, (New York: Wiley).
- MANSKI, C. (1993), "Dynamic Choice in Social Settings", *Journal of Econometrics*, 58, 121-136.
- McFADDEN, D. (1973), "Conditional Logit Analysis of Qualitative Choice Behavior", in P. Zarembka (ed.), *Frontiers in Econometrics*, (New York: Academic Press).
- McFADDEN, D. (1974), "The Measurement of Urban Travel Demand", *Journal of Public Economics*, 3, 303-328.
- McFADDEN, D. (1981), "Econometric Models of Probabilistic Choice", in C. Manski and C. McFadden (eds.), *Structural Analysis of Discrete Data with Econometric Applications*, (Cambridge: MIT Press).
- McFADDEN, D. (1984), "Qualitative Response Models", in Z. Griliches and M. Intriligator (eds.), *Handbook of Econometrics*, (Amsterdam: North Holland).
- McFADDEN, D. and RICHTER, M. (1970), "Revealed Stochastic Preference", *manuscript*, MIT.
- MILLER, R. (1984), "Job Matching and Occupational Choice", *Journal of Political Economy*, 92, 1086-1120.
- MILLER, R. (1997), "Estimating Models of Dynamic Optimization with Microeconomic Data", in H. Pesaran and P. Schmidt (eds.), *Handbook of Ap-*

- plied Econometrics, vol. II (Oxford: Blackwell)
- NILSEN O. and SCHIANTARELLI, F. (1996), "Zeroes and Lumps in Investment: empirical evidence on irreversibilities and non-convexities", manuscript.
- PAKES, A. (1986), "Patent as Options: Some Estimates of the Value of Holding European Patent Stocks", *Econometrica*, 54, 755-785.
- PAKES, A. (1994), "Dynamic Structural Models, Problems and Prospects: Mixed Continuous Discrete Controls and Market Interaction", in C. Sims (ed.), *Advances in Econometrics, Sixth World Congress*, vol. II.
- ROTA, P. (2000), "Dynamic Labour Demand with Lumpy and Kinked Adjustment Costs", *manuscript*.
- ROTA, P. (2000), "Empirical Representation of Firms' Employment Decisions by an (S,s) Rule", *manuscript*.
- RUST, J. (1987), "Optimal Replacement of GMC Bus Engines: an Empirical Model of Harold Zurcher", *Econometrica*, 55, 999-1033.
- RUST, J. (1988), "A Backward Solution Algorithm for Structural Estimation of Discrete Control Processes" *manuscript*.
- RUST, J. (1992), "Structural Estimation of Markov Decision Processes," in D. McFadden and R. Engle (eds.), *Handbook of Econometrics*, vol. 4 (Amsterdam: North Holland).
- RUST, J. (1994), "Estimation of Dynamic Structural Models: Problems and Prospects Part I: Discrete Decision Processes," in C. Sims (eds), *Proceedings of the 6th World Congress of the Econometric Society Barcelona*, Spain (Cambridge, U.K: Cambridge University Press).
- SCARF, (1960), "The Optimality of (s,S) Policies for the Dynamic Inventory Problem", *Proceedings of the 1st Stanford Symposium on Mathematical Methods in the Social Sciences*, Stanford University Press, Stanford, Cal.
- SLADE, M. (1998), "Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices", *Review of Economic Studies*, 65, 87-107.
- STOKEY, N. and LUCAS, R. (1989), *Recursive Methods in Economic Dynamics*, (Cambridge, Massachusetts: Harvard University Press).
- STURM, R. (1991), "A Structural Economic Model of Operating Cycle Management in European Nuclear Power Plants", *manuscript*.
- WHITE, H. (1980), "A Heteroscedasticity Consistent Covariance Matrix Estimator and a Direct Test for Heteroscedasticity", *Econometrica*, 48, 817-838.
- WOLPIN, K. (1987) "Estimating a Structural Search Model: The Transition from Schooling to Work", *Econometrica*, 55, 801-818.