

Abbring, Jaap H.; Heckman, James J.

Working Paper

Dynamic policy analysis

cemmap working paper, No. CWP05/08

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Abbring, Jaap H.; Heckman, James J. (2008) : Dynamic policy analysis, cemmap working paper, No. CWP05/08, Centre for Microdata Methods and Practice (cemmap), London, <https://doi.org/10.1920/wp.cem.2008.0508>

This Version is available at:

<https://hdl.handle.net/10419/79368>

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.

Dynamic policy analysis

Jaap H. Abbring
James J. Heckman

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP5/08

Dynamic Policy Analysis*

Jaap H. Abbring[†]
James J. Heckman[‡]

February 6, 2008

Abstract

This chapter studies the microeconomic treatment-effect and structural approaches to dynamic policy evaluation. First, we discuss a reduced-form approach based on a sequential randomization or dynamic matching assumption that is popular in biostatistics. We then discuss two complementary approaches for treatments that are single stopping times and that allow for non-trivial dynamic selection on unobservables. The first builds on continuous-time duration and event-history models. The second extends the discrete-time dynamic discrete-choice literature.

1 Introduction

The methods discussed in Parts 1 and 2 of this volume are useful for microeconomic policy evaluation. That field analyzes the effects of policy interventions on individual outcomes. Panel data facilitate the identification and estimation of such effects. Panel data are especially helpful in analyzing the individual dynamic consequences of policies and outcomes, which are mostly neglected in the vast cross-sectional literature on this topic. Not surprisingly, panel-data methods are becoming more widely used in the microeconomic policy evaluation literature. In this chapter we critically review recently developed methods and their applications.

The outline of the chapter is as follows. Section 2 presents the policy evaluation problem and discusses the treatment-effect approach to policy evaluation. It establishes the notation used in the rest of this chapter. Section 3 reviews an approach to the analysis of dynamic treatment effects based on a sequential randomization assumption that is popular in biostatistics (Gill and Robins, 2001; Lok, 2007; Robins, 1997) and has been applied in economics (see Fitzenberger, Osikominu, and Völter, 2006, and Lechner and Miquel, 2002). This is a dynamic version of matching. We relate the assumptions justifying this approach to the assumptions underlying the econometric

*Paper forthcoming as Chapter 24 in L. Matyas and P. Sevestre (2008), editors, *The Econometrics of Panel Data*, third edition, Kluwer, Dordrecht.

[†]Department of Economics, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, and Tinbergen Institute. Email: jabbring@econ.vu.nl.

[‡]Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago IL 60637, USA; American Bar Foundation; and Geary Institute, University College Dublin. Email: jjh@uchicago.edu.

dynamic discrete-choice literature based on Rust’s (1987) conditional-independence condition which, as discussed in section 5.5 below, is frequently invoked in the structural econometrics literature. We note the limitations of the dynamic matching treatment-effect approach in accounting for dynamic information accumulation. In sections 4 and 5, we discuss two econometric approaches for the analysis of treatment times that allow for non-trivial dynamic selection on unobservables. Section 4 discusses the continuous-time event-history approach to policy evaluation (Abbring and Van den Berg, 2003b, 2005) and Abbring (2008). Section 5 introduces an approach developed by Heckman and Navarro (2007) that builds on and extends the discrete-time dynamic discrete-choice literature. Like the analysis of Abbring and Van den Berg, it does not rely on the conditional-independence assumptions used in dynamic matching. The two complementary approaches surveyed in this chapter span the existing econometric literature on dynamic treatment effects.

2 Policy Evaluation and Treatment Effects

2.1 The Evaluation Problem

We introduce some key ideas and set up the notation for this chapter by reviewing the static policy evaluation problem discussed in, e.g., Heckman and Vytlacil (2007a). Let Ω be the set of agent types. It is the sample space of a probability space $(\Omega, \mathcal{I}, \mathbb{P})$, and all choices and outcomes are random variables defined on this probability space. Each agent type $\omega \in \Omega$ represents a single agent in a particular state of nature. We could distinguish variation between agents from within-agent randomness by taking $\Omega = J \times \tilde{\Omega}$, with J the set of agents and $\tilde{\Omega}$ the set of possible states of nature. However, we do not make this distinction explicit in this chapter, and often simply refer to agents instead of agent types.¹

Consider a policy that targets the allocation of each agent in Ω to a single treatment from a set \mathcal{S} . In the most basic binary version, $\mathcal{S} = \{0, 1\}$, where “1” represents “treatment”, such as a training program, and “0” some baseline, “control” program. Alternatively, \mathcal{S} could take a continuum of values, e.g., $\mathbb{R}_+ = [0, \infty)$, representing, e.g., unemployment benefit levels, or duration of time in a program.

A policy $p = (a, \tau) \in \mathcal{A} \times \mathcal{T} \equiv \mathcal{P}$ consists of a planner’s rule $a : \Omega \rightarrow \mathcal{B}$ for allocating constraints and incentives to agents, and a rule $\tau : \Omega \times \mathcal{A} \rightarrow \mathcal{S}$ that generates agent treatment choices for a given constraint allocation a . This framework allows agent ω ’s treatment choice to depend both on the constraint assignment mechanism a — in particular, the distribution of the constraints in the population— and on the constraints $a(\omega) \in \mathcal{B}$ assigned to agent ω .

The randomness in the planner’s constraint assignment a may reflect heterogeneity of agents as observed by the planner, but it may also be due to explicit randomization. For example, consider profiling on background characteristics of potential participants in the assignment a to treatment eligibility. If the planner observes some background characteristics on individuals in the population of interest, she could choose eligibility status to be a deterministic function of those characteristics

¹For example, we could have $\Omega = [0, 1]$ indexing the population of agents, with \mathbb{P} being Lebesgue measure on $[0, 1]$. Alternatively, we could take $\Omega = [0, 1] \times \tilde{\Omega}$ and have $[0, 1]$ represent the population of agents and $\tilde{\Omega}$ states of nature.

and, possibly, some other random variable under her control by randomization. This includes the special case in which the planner randomizes persons into eligibility. We denote the information set generated by the variables observed by the planner when she assigns constraints, including those generated through deliberate randomization, by \mathcal{I}_P .² The planner's information set \mathcal{I}_P determines how precisely she can target each agent's ω when assigning constraints. The variables in the information set fully determine the constraints assignment a .

Subsequent to the planner's constraints assignment a , each agent ω chooses treatment $\tau(\omega, a)$. We assume that agents know the constraint assignment mechanism a in place. However, agents do not directly observe their types ω , but only observe realizations $I_A(\omega)$ of some random variables I_A . For given $a \in \mathcal{A}$, agent ω 's treatment choice $\tau(\omega, a)$ can only depend on ω through his observations $I_A(\omega)$. Typically, $I_A(\omega)$ includes the variables used by the planner in determining $a(\omega)$, so that agents know the constraints that they are facing. Other components of $I_A(\omega)$ may be determinants of preferences and outcomes. Variation in $I_A(\omega)$ across ω may thus reflect preference heterogeneity, heterogeneity in the assigned constraints, and heterogeneity in outcome predictors. We use \mathcal{I}_A to denote the information set generated by I_A .³ An agent's information set \mathcal{I}_A determines how precisely the agent can tailor his treatment choice to his type ω . For expositional convenience, we assume that agents know more when choosing treatment than what the planner knows when assigning constraints, so that $\mathcal{I}_A \supseteq \mathcal{I}_P$. One consequence is that agents observe the constraints $a(\omega)$ assigned to them, as previously discussed. In turn, the econometrician may not have access to all of the information that is used by the agents when they choose treatment.⁴ In this case, $\mathcal{I}_A \not\subseteq \mathcal{I}_E$, where \mathcal{I}_E denotes the econometrician's information set.

We define $s_p(\omega)$ as the treatment selected by agent ω under policy p . With $p = (a, \tau)$, we have that $s_p(\omega) = \tau(\omega, a)$. The random variable $s_p : \Omega \rightarrow \mathcal{S}$ represents the allocation of agents to treatments implied by policy p .⁵ Randomness in this allocation reflects both heterogeneity in the planner's assignment of constraints and the agents' heterogeneous responses to this assignment. One extreme case arises if the planner assigns agents to treatment groups and agents perfectly comply, so that $\mathcal{B} = \mathcal{S}$ and $s_p(\omega) = \tau(\omega, a) = a(\omega)$ for all $\omega \in \Omega$. In this case, all variation of s_p is due to heterogeneity in the constraints $a(\omega)$ across agents ω . At the other extreme, agents do not respond at all to the incentives assigned by mechanisms in \mathcal{A} , and $\tau(a, \omega) = \tau(a', \omega)$ for all $a, a' \in \mathcal{A}$ and $\omega \in \Omega$. In general, there are policies that have a nontrivial (that is, nondegenerate) constraint

²Formally, \mathcal{I}_P is a sub- σ -algebra of \mathcal{I} and a is assumed to be \mathcal{I}_P -measurable.

³Formally, \mathcal{I}_A is a sub- σ -algebra of \mathcal{I} — the σ -algebra generated by I_A — and $\omega \in \Omega \mapsto \tau(\omega, a) \in \mathcal{S}$ should be \mathcal{I}_A -measurable for all $a \in \mathcal{A}$. The possibility that different agents have different information sets is allowed for because a distinction between agents and states of nature is implicit. As suggested in the introduction to this section, we can make it explicit by distinguishing a set J of agents and a set $\tilde{\Omega}$ of states of nature and writing $\Omega = J \times \tilde{\Omega}$. For expositional convenience, let J be finite. We can model the case that agents observe their identity j by assuming that the random variable J_A on Ω that reveals their identity, that is $J_A(j, \tilde{\omega}) = j$, is in their information set \mathcal{I}_A . If agents, in addition, observe some other random variable V on Ω , then the information set \mathcal{I}_A generated by (J_A, V) can be interpreted as providing each agent $j \in J$ with perfect information about his identity j and with the agent- j -specific information about the state of nature $\tilde{\omega}$ encoded in the random variable $\tilde{\omega} \mapsto V(j, \tilde{\omega})$ on $\tilde{\Omega}$.

⁴See the discussion by Heckman and Vytlačil (2007b, sections 2 and 9).

⁵Formally, $\{s_p\}_{p \in \mathcal{A} \times \mathcal{T}}$ is a stochastic process indexed by p .

assignment a , where at least some agents respond to the assigned constraints a in their treatment choice, $\tau(a, \omega) \neq \tau(a', \omega)$ for some $a, a' \in \mathcal{A}$ and $\omega \in \Omega$.

We seek to evaluate a policy p in terms of some outcome Y_p , for example, earnings. For each $p \in \mathcal{P}$, Y_p is a random variable defined on the population Ω . The evaluation can focus on objective outcomes Y_p , on the subjective valuation $R(Y_p)$ of Y_p by the planner or the agents, or on both types of outcomes. The evaluation can be performed relative to a variety of information sets reflecting different actors (the agent, the planner and the econometrician) and the arrival of information in different time periods. Thus, the randomness of Y_p may represent both (*ex ante*) heterogeneity among agents known to the planner when constraints are assigned (that is, variables in \mathcal{I}_P) and/or heterogeneity known to the agents when they choose treatment (that is, information in \mathcal{I}_A), as well as (*ex post*) shocks that are not foreseen by the policy maker or by the agents. An information-feasible (*ex ante*) policy evaluation by the planner would be based on some criterion using the distribution of Y_p conditional on \mathcal{I}_P . The econometrician can assist the planner in computing this evaluation if the planner shares her *ex ante* information and $\mathcal{I}_P \subseteq \mathcal{I}_E$.

Suppose that we have data on outcomes Y_{p_0} under policy p_0 with corresponding treatment assignment s_{p_0} . Consider an intervention that changes the policy from the actual p_0 to some counterfactual p' with associated treatments $s_{p'}$ and outcomes $Y_{p'}$. This could involve a change in the planner's constraint assignment from a_0 to a' for given $\tau_0 = \tau'$, a change in the agent choice rule from τ_0 to τ' for given $a_0 = a'$, or both.

The policy evaluation problem involves contrasting $Y_{p'}$ and Y_{p_0} or functions of these outcomes. For example, if the outcome of interest is mean earnings, we might be interested in some weighted average of $E[Y_{p'} - Y_{p_0} \mid \mathcal{I}_P]$, such as $E[Y_{p'} - Y_{p_0}]$. The special case where $\mathcal{S} = \{0, 1\}$ and $s_{p'} = a' = 0$ generates the effect of abolishing the program. Implementing such a policy requires that the planner be able to induce all agents into the control group by assigning constraints $a' = 0$. In particular, this assumes that there are no substitute programs available to agents that are outside the planner's control (Heckman and Vytlačil, 2007b, section 10).

For notational convenience, write $S = s_{p_0}$ for treatment assignment under the actual policy p_0 in place. Cross-sectional micro data typically provide a random sample from the joint distribution of (Y_{p_0}, S) .⁶ Clearly, without further assumptions, such data do not identify the effects of the policy shift from p_0 to p' . This identification problem becomes even more difficult if we do not seek to compare the counterfactual policy p' with the actual policy p_0 , but rather with another counterfactual policy p'' that also has never been observed. A leading example is the binary case in which $0 < \Pr(S = 1) < 1$, but we seek to know the effects of $s_{p'} = 0$ (universal nonparticipation) and $s_{p''} = 1$ (universal treatment), where neither policy has ever been observed in place.

Panel data can help to evaluate the type of static policies discussed so far, if interpreted as short-run or even one-shot policies. Suppose that we have data on outcomes in two periods in which two different policies were in place. In a world in which outcomes in any period are not affected by the policy or outcomes in the other period, such data are directly informative on the contrast

⁶Notice that a random sample of outcomes under a policy may entail nonrandom selection of treatments as individual agents select individual treatments given τ and the constraints they face assigned by a .

between outcomes under both policies.

The standard microeconomic approach to the policy evaluation problem assumes that the (subjective and objective) outcomes for any individual agent are the same across all policy regimes for any particular treatment assigned to the individual (see, e.g., Heckman, LaLonde, and Smith, 1999). Heckman and Vytlacil (2007a) present a detailed account of the policy-invariance assumptions that justify this practice. They simplify the task of evaluating policy p to determining (i) the assignment s_p of treatments under policy p and (ii) treatment effects for individual outcomes. Even within this simplified framework, there are still two difficult, and distinct, problems in identifying treatment effects on individual outcomes:

- (A) *The Evaluation Problem: that we observe an agent in one treatment state and seek to determine that agent's outcomes in another state; and*
- (B) *The Selection Problem: that the distributions of outcomes for the agents we observe in a given treatment state are not the marginal population distributions that would be observed if agents were randomly assigned to the state.*

The assignment mechanism s_p of treatments under counterfactual policies p is straightforward in the case where the planner assigns agents to treatment groups and agents fully comply, so that $s_p = a$. More generally, an explicit model of agent treatment choices is needed to derive s_p for counterfactual policies p . An explicit model of agent treatment choices can also be helpful in addressing the selection problem, and in identifying agent subjective valuations of outcomes. We now formalize the notation for the treatment-effect approach that we will use in this section using the potential-outcome framework of Neyman (1923), Roy (1951), Quandt (1958, 1972) and Rubin (1974).⁷

2.2 The Treatment-Effect Approach

For each agent $\omega \in \Omega$, let $y(s, X(\omega), U(\omega))$ be the potential outcome when the agent is assigned to treatment $s \in \mathcal{S}$. Here, X and U are covariates that are not causally affected by the treatment or the outcomes.^{8,9} In the language of Kalbfleisch and Prentice (1980) and Leamer (1985), we say that such covariates are “external” to the causal model. X is observed by the econometrician (that is, in \mathcal{I}_E) and U is not.

Recall that s_p is the assignment of agents to treatments under policy p . For all policies p that we consider, the outcome Y_p is linked to the potential outcomes by the consistency condition

⁷See Heckman, LaLonde, and Smith (1999), Heckman and Vytlacil (2007a), and Thurstone (1930) for results in econometrics and extensive reviews of the econometric literature.

⁸This “no feedback” condition requires that X and U are the same fixing treatment to s for all s . See Haavelmo (1943), Pearl (2000), or the discussion in Heckman and Vytlacil (2007a,b).

⁹Note that this framework is rich enough to capture the case in which potential outcomes depend on treatment-specific unobservables as in section 5, because these can be simply stacked in U and subsequently selected by y . For example, in the case where $\mathcal{S} = \{0, 1\}$ we can write $y(s, X, (U_0, U_1)) = sy_1(X, U_1) + (1 - s)y_0(X, U_0)$ for some y_0 and y_1 . A specification without treatment-dependent unobservables is more tractable in the case of continuous treatments in section 3 and, in particular, continuous treatment times in section 4.

$Y_p = y(s_p, X, U)$. This condition follows from the policy-invariance assumptions. It embodies the assumption that an agent's outcome only depends on the treatment assigned to the agent and not separately on the mechanism used to assign treatments. This excludes (strategic) interactions between agents and equilibrium effects of the policy.¹⁰ It ensures that we can specify individual outcomes y from participating in programs in \mathcal{S} independently of the policy p and treatment assignment s_p . Economists say that y is autonomous, or structurally invariant with respect to the policy environment (see Frisch, 1938; Hurwicz, 1962; and Heckman and Vytlačil, 2007a).¹¹

We illustrate the treatment-effect approach with a basic example. Consider the evaluation of an intervention that changes the policy from p_0 to p' in terms of its mean effect $E[Y_{p'} - Y_{p_0}]$ on outcomes. For expositional convenience, let treatment be binary: $\mathcal{S} = \{0, 1\}$. Suppose that we have a cross-sectional sample from the joint distribution of (Y_{p_0}, S, X) . Assume that treatment assignment under both the actual policy p_0 and the alternative policy p' is randomized, that is, both S and $s_{p'}$ are independent of the determinants (X, U) of the potential outcomes. Then, because of the policy-invariance conditions,

$$E[Y_{p'} - Y_{p_0}] = E[y(1, X, U) - y(0, X, U)] [\Pr(s_{p'} = 1) - \Pr(S = 1)] .$$

The mean effect of the intervention on outcomes equals the ‘average treatment effect’ $E[y(1, X, U) - y(0, X, U)]$ times the net increase in the assignment to treatment 1. The policy evaluation problem boils down to identifying the average treatment effect, the distribution of the actual treatment assignment S , and the distribution of treatment assignment $s_{p'}$ under the alternative policy p' . Under the assumption of randomized assignment, and provided that $0 < \Pr(S = 1) < 1$, the average treatment effect is identified as $E[Y_{p_0}|S = 1] - E[Y_{p_0}|S = 0]$. The distribution of S is identified directly from the data. The distribution of $s_{p'}$ is often known, as in the case of universal nonparticipation ($s_{p'} = 0$) or universal treatment ($s_{p'} = 1$). Otherwise, it needs to be identified using a model of treatment choice.

Heckman and Vytlačil (2007a,b) review more general evaluation problems and econometric methods that do not rely on randomized assignment, such as the methods of matching and instrumental variables. Clearly, panel data, combined with stationarity assumptions, can help in addressing the selection problem in the evaluation of static policies. We will not dwell on this application of panel data to the evaluation of static policies, but now turn to the dynamic policy evaluation problem.

2.3 Dynamic Policy Evaluation

Interventions often have consequences that span over many periods. Policy interventions at different points in time can be expected to affect not only current outcomes, but also outcomes at other points in time. The same policy implemented at different time periods may have different consequences. Moreover, policy assignment rules often have non-trivial dynamics. The assignment of programs at

¹⁰See Pearl (2000), Heckman (2005), and Heckman and Vytlačil (2007a).

¹¹See also Aldrich (1989) and Hendry and Morgan (1995). Rubin's (1986) stable-unit-treatment-value assumption is a version of the classical invariance assumptions of econometrics (see Abbring, 2003, and Heckman and Vytlačil, 2007a, for discussion of this point).

any point in time can be contingent on the available data on past program participation, intermediate outcomes and covariates.

The dynamic policy evaluation problem can be formalized in a fashion similar to the way we formalized the static problem in subsection 2.1. In this subsection, we analyze a discrete-time finite-horizon model. We consider continuous-time models in section 4. The possible treatment assignment times are $1, \dots, \bar{T}$. We do not restrict the set \mathcal{S} of treatments. We allow the same treatment to be assigned on multiple occasions. In general, the set of available treatments at each time t may depend on time t and on the history of treatments, outcomes, and covariates. For expositional convenience, we will only make this explicit in sections 4 and 5, where we focus on the timing of a single treatment.

We define a dynamic policy $p = (a, \tau) \in \mathcal{A} \times \mathcal{T} \equiv \mathcal{P}$ as a dynamic constraint assignment rule $a = \{a_t\}_{t=1}^{\bar{T}}$ with a dynamic treatment choice rule $\tau = \{\tau_t\}_{t=1}^{\bar{T}}$. At each time t , the planner assigns constraints $a_t(\omega)$ to each agent $\omega \in \Omega$, using information in the time- t policy- p information set $\mathcal{I}_P(t, p) \subseteq \mathcal{I}$. The planner's information set $\mathcal{I}_P(t, p)$ could be based on covariates and random variables under the planner's control, as well as past choices and realized outcomes. We denote the sequence of planner's information sets by $\mathcal{I}_P(p) = \{\mathcal{I}_P(t, p)\}_{t=1}^{\bar{T}}$. We assume that the planner does not forget any information she once had, so that her information improves over time and $\mathcal{I}_P(t, p) \subseteq \mathcal{I}_P(t+1, p)$ for all t .¹²

Each agent ω chooses treatment $\tau_t(\omega, a)$ given their information about ω at time t under policy p and given the constraint assignment mechanism $a \in \mathcal{A}$ in place. We assume that agents know the constraint assignment mechanism a in place. At time t , under policy p , agents infer their information about their type ω from random variables $I_A(t, p)$ that may include preference components and determinants of constraints and future outcomes. $\mathcal{I}_A(t, p)$ denotes the time- t policy- p information set generated by $I_A(t, p)$ and $\mathcal{I}_A(p) = \{\mathcal{I}_A(t, p)\}_{t=1}^{\bar{T}}$. We assume that agents are increasingly informed as time goes by, so that $\mathcal{I}_A(t, p) \subseteq \mathcal{I}_A(t+1, p)$.¹³ For expositional convenience, we also assume that agents know more than the planner at each time t , so that $\mathcal{I}_P(t, p) \subseteq \mathcal{I}_A(t, p)$.¹⁴ Because all determinants of past and current constraints are in the planner's information set $\mathcal{I}_P(t, p)$, this implies that agents observe $(a_1(\omega), \dots, a_t(\omega))$ at time t . Usually, they do not observe all determinants of their future constraints $(a_{t+1}(\omega), \dots, a_{\bar{T}}(\omega))$.¹⁵ Thus, the treatment choices of the agents may be contingent on past and current constraints, their preferences, and on their predictions of future outcomes and constraints given their information $\mathcal{I}_A(t, p)$ and given the constraint assignment mechanism a in place.

Extending the notation for the static case, we denote the assignment of agents to treatment τ_t at time t implied by a policy p by the random variable $s_p(t)$ defined so that $s_p(\omega, t) = \tau_t(\omega, a)$. We use the shorthand s_p^t for the vector $(s_p(1), \dots, s_p(t))$ of treatments assigned up to and including time t

¹²Formally, the information $\mathcal{I}_P(p)$ that accumulates for the planner under policy p is a filtration in \mathcal{I} , and a is a stochastic process that is adapted to $\mathcal{I}_P(p)$.

¹³Formally, the information $\mathcal{I}_A(p)$ that accumulates for the agents is a filtration in \mathcal{I} .

¹⁴If agents are strictly better informed, and $\mathcal{I}_P(t, p) \subset \mathcal{I}_A(t, p)$, it is unlikely that the planner catches up and learns the agent's information with a delay (e.g., $\mathcal{I}_A(t, p) \subseteq \mathcal{I}_P(t+1, p)$) unless agent's choices and outcomes reveal all their private information.

¹⁵Formally, a_1, \dots, a_t are $\mathcal{I}_A(t, p)$ -measurable, but $a_{t+1}, \dots, a_{\bar{T}}$ are not.

under policy p , and write $s_p = s_p^{\bar{T}}$. The assumptions made so far about the arrival of information imply that treatment assignment $s_p(t)$ can only depend on the information $\mathcal{I}_A(t, p)$ available to agents at time t .¹⁶

Because past outcomes typically depend on the policy p , the planner's information $\mathcal{I}_P(p)$ and the agents' information $\mathcal{I}_A(p)$ will generally depend on p as well. In the treatment-effect framework that we develop in the next section, at each time t different policies may have selected different elements in the set of potential outcomes in the past. The different elements reveal different aspects of the unobservables underlying past and future outcomes. We will make assumptions that limit the dependence of information sets on policies in the context of the treatment-effect approach developed in the next section.

Objective outcomes associated with policies p are expressed as a vector of time-specific outcomes $Y_p = (Y_p(1), \dots, Y_p(\bar{T}))$. The components of this vector may also be vectors. We denote the outcomes from time 1 to time t under policy p by $Y_p^t = (Y_p(1), \dots, Y_p(t))$. We analyze both subjective and objective evaluations of policies in section 5, where we consider more explicit economic models. Analogous to our analysis of the static case, we cannot learn about the outcomes $Y_{p'}$ that would arise under a counterfactual policy p' from data on outcomes Y_{p_0} and treatments $s_{p_0} = S$ under a policy $p_0 \neq p'$ without imposing further structure on the problem.¹⁷ We follow the approach explicated for the static case and assume policy invariance of individual outcomes under a given treatment. This reduces the evaluation of a dynamic policy p to identifying (i) the dynamic assignment s_p of treatments under policy p and (ii) the dynamic treatment effects on individual outcomes. We focus our discussion on the fundamental evaluation problem and the selection problem that haunt inference about treatment effects. In the remainder of the section, we review alternative approaches to identifying dynamic treatment effects, and some approaches to modeling dynamic treatment choice. We first analyze methods recently developed in statistics.

3 Dynamic Treatment Effects and Sequential Randomization

In a series of papers, Robins extends the static Neyman-Roy-Rubin model based on selection on observables to a dynamic setting (see, e.g., Robins, 1997, and the references therein). He does not consider agent choice or subjective evaluations. Here, we review his extension, discuss its relationship to dynamic choice models in econometrics, and assess its merits as a framework for economic policy analysis. We follow the exposition of Gill and Robins (2001), but add some additional structure to their basic framework to explicate the connection of their approach to the dynamic approach pursued in econometrics.

¹⁶Formally, $\{s_p(t)\}_{t=1}^{\bar{T}}$ is a stochastic process that is adapted to $\mathcal{I}_A(p)$.

¹⁷If outcomes under different policy regimes are informative about the same technology and preferences, for example, then the analyst and the agent could learn about the ingredients that produce counterfactual outcomes in all outcome states.

3.1 Dynamic Treatment Effects

3.1.1 Dynamic Treatment and Dynamic Outcomes

To simplify the exposition, suppose that \mathcal{S} is a finite discrete set.¹⁸ Recall that, at each time t and for given p , treatment assignment $s_p(t)$ is a random variable that only depends on the agent's information $\mathcal{I}_A(t, p)$, which includes personal knowledge of preferences and determinants of constraints and outcomes. To make this dependence explicit, suppose that external covariates Z , observed by the econometrician (that is, variables in \mathcal{I}_E), and unobserved external covariates V_1 that affect treatment assignment are revealed to the agents at time 1. Then, at the start of each period $t \geq 2$, past outcomes $Y_p(t-1)$ corresponding to the outcomes realized under treatment assignment s_p and external unobserved covariates V_t enter the agent's information set.¹⁹ In this notation, $\mathcal{I}_A(1, p)$ is the information $\sigma(Z, V_1)$ conveyed to the agent by (Z, V_1) and, for $t \geq 2$, $\mathcal{I}_A(t, p) = \sigma(Y_p^{t-1}, Z, V^t)$, with $V^t = (V_1, \dots, V_t)$. In the notation of the previous subsection, $I_A(1, p) = (Z, V_1)$ and, for $t \geq 2$, $I_A(t, p) = (Y_p^{t-1}, Z, V^t)$. Among the elements of $I_A(t, p)$ are the determinants of the constraints faced by the agent up to t , which may or may not be observed by the econometrician.

We attach *ex post* potential outcomes $Y(t, s) = y_t(s, X, U_t)$, $t = 1, \dots, \bar{T}$, to each treatment sequence $s = (s(1), \dots, s(\bar{T}))$. Here, X is a vector of observed (by the econometrician) external covariates and U_t , $t = 1, \dots, \bar{T}$, are vectors of unobserved external covariates. Some components of X and U_t may be in agent information sets. We denote $Y^t(s) = (Y(1, s), \dots, Y(t, s))$, $Y(s) = Y^{\bar{T}}(s)$, and $U = (U_1, \dots, U_{\bar{T}})$. As in the static case, potential outcomes y are assumed to be invariant across policies p , which ensures that $Y_p(t) = y_t(s_p, X, U_t)$. In the remainder of this section, we keep the dependence of outcomes on observed covariates X implicit and suppress all conditioning on X .

We assume no causal dependence of outcomes on future treatment:²⁰

(NA) For all $t \geq 1$, $Y(t, s) = Y(t, s')$ for all s, s' such that $s^t = (s')^t$,

where $s^t = (s(1), \dots, s(t))$ and $(s')^t = (s'(1), \dots, s'(t))$. Abbring and Van den Berg (2003b) and Abbring (2003) define this as a “no-anticipation” condition. It requires that outcomes at time t (and before) be the same across policies that allocate the same treatment up to and including t , even if they allocate different treatments after t . In the structural econometric models discussed in sections 3.2.2 and 5 below, this condition is trivially satisfied if all state variables relevant to outcomes at time t are included as inputs in the outcome equations $Y(t, s) = y_t(s, U_t)$, $t = 1, \dots, \bar{T}$.

Because Z and V_1 are assumed to be externally determined, and therefore not affected by the policy p , the initial agent information set $\mathcal{I}_A(1, p) = \sigma(Z, V_1)$ does not depend on p . Agent ω has the same initial data $(Z(\omega), V_1(\omega))$ about his type ω under all policies p . Thus, $\mathcal{I}_A(1, p) = \mathcal{I}_A(1, p')$ is a natural benchmark information set for an *ex ante* comparison of outcomes at time 1 among different policies. For $t \geq 2$, (NA) implies that actual outcomes up to time $t-1$ are

¹⁸All of the results presented in this subsection extend to the case of continuous treatments. We will give references to the appropriate literature in subsequent footnotes.

¹⁹Note that any observed covariates that are dynamically revealed to the agents can be subsumed in the outcomes.

²⁰For statistical inference from data on the distribution of (Y_{p_0}, S, Z) , these equalities only need to hold on events $\{\omega \in \Omega : S^t(\omega) = s^t\}$, $t \geq 1$, respectively.

equal between policies p and p' , $Y_p^{t-1} = Y_{p'}^{t-1}$, if the treatment histories coincide up to time $t - 1$ so that $s_p^{t-1} = s_{p'}^{t-1}$. Together with the assumption that Z and V^t are externally determined, it follows that agents have the same time- t information set structure about ω under policies p and p' , $\mathcal{I}_A(t, p) = \sigma(Y_p^{t-1}, Z, V^t) = \sigma(Y_{p'}^{t-1}, Z, V^t) = \mathcal{I}_A(t, p')$, if $s_p^{t-1} = s_{p'}^{t-1}$.^{21,22} In this context, $\mathcal{I}_A(t, p) = \mathcal{I}_A(t, p')$ is a natural information set for an *ex ante* comparison of outcomes from time t onwards between any two policies p and p' such that $s_p^{t-1} = s_{p'}^{t-1}$.

With this structure on the agent information sets in hand, it is instructive to review the separate roles in determining treatment choice of information about ω and knowledge about the constraint assignment rule a . First, agent ω 's time- t treatment choice $s_p(\omega, t) = \tau_t(\omega, a)$ may depend on distributional properties of a , for example the share of agents assigned to particular treatment sequences, and on the past and current constraints $(a_1(\omega), \dots, a_t(\omega))$ that were actually assigned to him. We have assumed both to be known to the agent. Both may differ between policies, even if the agent information about ω is fixed across the policies. Second, agent ω 's time- t treatment choice may depend on agent ω 's predictions of future constraints and outcomes. A forward-looking agent ω will use observations of his covariates $Z(\omega)$ and $V^t(\omega)$ and past outcomes $Y_p^{t-1}(\omega)$ to infer his type ω and subsequently predict future external determinants $(U_t(\omega), \dots, U_{\bar{T}}(\omega))$ of his outcomes and $(V_{t+1}(\omega), \dots, V_{\bar{T}}(\omega))$ of his constraints and treatments. In turn, this information updating allows agent ω to predict his future potential outcomes $(Y(t, s, \omega), \dots, Y(\bar{T}, s, \omega))$ and, for a given policy regime p , his future constraints $(a_{t+1}(\omega), \dots, a_{\bar{T}}(\omega))$, treatments $(s_p(t+1, \omega), \dots, s_p(\bar{T}, \omega))$, and realized outcomes $(Y_p(t, \omega), \dots, Y_p(\bar{T}, \omega))$. Under different policies, the agent may gather different information on his type ω and therefore come up with different predictions of the external determinants of his future potential outcomes and constraints. In addition, even if the agent has the same time- t predictions of the external determinants of future constraints and potential outcomes, he may translate these into different predictions of future constraints and outcomes under different policies.

Assumption (NA) requires that current potential outcomes are not affected by future treatment. Justifying this assumption requires specification of agent information about future treatment and agent behavior in response to that information. Such an interpretation requires that we formalize how information accumulates for agents across treatment sequences s and s' such that $s^t = (s')^t$ and $(s_{t+1}, \dots, s_{\bar{T}}) \neq (s'_{t+1}, \dots, s'_{\bar{T}})$. To this end, consider policies p and p' such that $s_p = s$ and $s_{p'} = s'$. These policies produce the same treatment assignment up to time t , but are different in the future. We have previously shown that, even though the time- t agent information about ω is the same under both policies, $\mathcal{I}_A(t, p) = \mathcal{I}_A(t, p')$, agents may have different predictions of future constraints, treatments and outcomes because the policies may differ in the future and agents know

²¹If $s_p^{t-1}(\omega) = s_{p'}^{t-1}(\omega)$ only holds for ω in some subset $\Omega_{t-1} \subset \Omega$ of agents, then $Y_p^{t-1}(\omega) = Y_{p'}^{t-1}(\omega)$ only for $\omega \in \Omega_{t-1}$, and information coincides between p and p' only for agents in Ω_{t-1} . Formally, let Ω_{t-1} be the set $\{\omega \in \Omega : s_p^{t-1}(\omega) = s_{p'}^{t-1}(\omega)\}$ of agents that share the same treatment up to and including time $t - 1$. Then, Ω_{t-1} is in the agent's information set under both policies, $\Omega_{t-1} \in \mathcal{I}_A(t, p) \cap \mathcal{I}_A(t, p')$. Moreover, the partitioning of Ω_{t-1} implied by $\mathcal{I}_A(t, p)$ and $\mathcal{I}_A(t, p')$ is the same. To see this, note that the collections of all sets in, respectively, $\mathcal{I}_A(t, p)$ and $\mathcal{I}_A(t, p')$ that are weakly included in Ω_{t-1} are identical σ -algebras on Ω_{t-1} .

²²Notice that the realizations of the random variables $Y_{p'}^{t-1}$, Z , V^t may differ among agents.

this. The policy-invariance conditions ensure that time- t potential outcomes are nevertheless the same under each policy. This requires that potential outcomes be determined externally, and are not affected by agent actions in response to different predictions of future constraints, treatments and outcomes. This assumption rules out investment responses to alternative policies that affect potential outcomes.

In general, different policies in \mathcal{P} will produce different predictions of future constraints, treatment and outcomes. In the dynamic treatment-effect framework, this may affect outcomes indirectly through agent treatment choices. If potential outcomes are directly affected by agent's forward-looking decisions, then the invariance conditions underlying the treatment-effect framework will be violated. Section 3.3 illustrates this issue, and the no-anticipation condition, with some examples.

3.1.2 Identification of Treatment Effects

Suppose that the econometrician has data that allows her to estimate the joint distribution of (Y_{p_0}, S, Z) of outcomes, treatments and covariates under some policy p_0 , where $S = s_{p_0}$. These data are not enough to identify dynamic treatment effects.

To secure identification, Gill and Robins (2001) invoke a dynamic version of the matching assumption (conditional independence) which relies on sequential randomization:²³

(M-1) *For all treatment sequences s and all t ,*

$$S(t) \perp\!\!\!\perp (Y(t, s), \dots, Y(\bar{T}, s)) \mid (Y_{p_0}^{t-1}, S^{t-1} = s^{t-1}, Z),$$

where the conditioning set $(Y_{p_0}^0, S^0 = s^0, Z)$ for $t = 1$ is Z .

Equivalently, $S(t) \perp\!\!\!\perp (U_t, \dots, U_{\bar{T}}) \mid (Y_{p_0}^{t-1}, S^{t-1}, Z)$ for all t without further restricting the data. Sequential randomization allows the $Y_{p_0}(t)$ to be “dynamic confounders”—variables that are affected by past treatment and that affect future treatment assignment.

The sequence of conditioning information sets appearing in the sequential randomization assumption, $\mathcal{I}_E(1) = \sigma(Z)$ and, for $t \geq 2$, $\mathcal{I}_E(t) = \sigma(Y_{p_0}^{t-1}, S^{t-1}, Z)$, is a filtration \mathcal{I}_E of the econometrician's information set $\sigma(Y_{p_0}, S, Z)$. Note that $\mathcal{I}_E(t) \subseteq \mathcal{I}_A(t, p_0)$ for each t . If treatment assignment is based on strictly more information than \mathcal{I}_E , so that agents know strictly more than the econometrician and act on their superior information, (M-1) is likely to fail if that extra information also affects outcomes. Heckman and Vytlačil (2007b) make this point in a static setting.

Together with the no-anticipation condition (NA), which is a condition on outcomes and distinct from (M-1), the dynamic potential-outcome model set up so far is a natural dynamic extension of the Neyman-Roy-Rubin model for a static (stratified) randomized experiment.

Under assumption (M-1) that the actual treatment assignment S is sequentially randomized, we can sequentially identify the causal effects of treatment from the distribution of the data (Y_{p_0}, S, Z) and construct the distribution of the potential outcomes $Y(s)$ for any treatment sequence s in the support of S .

²³Formally, we need to restrict attention to sequences s in the support of S . Throughout this section, we will assume this and related support conditions hold.

Consider the case in which all variables are discrete. No-anticipation condition (NA) ensures that potential outcomes for a treatment sequence s equal actual (under policy p_0) outcomes up to time $t - 1$ for agents with treatment history s^{t-1} up to time $t - 1$. Formally, $Y^{t-1}(s) = Y_{p_0}^{t-1}$ on the set $\{S^{t-1} = s^{t-1}\}$. Using this, sequential randomization assumption (M-1) can be rephrased in terms of potential outcomes: for all s and t ,

$$S(t) \perp\!\!\!\perp (Y(t, s), \dots, Y(\bar{T}, s)) \mid (Y^{t-1}(s), S^{t-1} = s^{t-1}, Z).$$

In turn, this implies that, for all s and t ,

$$\Pr(Y(t, s) = y(t) \mid Y^{t-1}(s) = y^{t-1}, S^t = s^t, Z) = \Pr(Y(t, s) = y(t) \mid Y^{t-1}(s) = y^{t-1}, Z), \quad (1)$$

where $y^{t-1} = (y(1), \dots, y(t-1))$ and $y = y^{\bar{T}}$. From Bayes' rule and (1), it follows that

$$\begin{aligned} & \Pr(Y(s) = y \mid Z) \\ &= \Pr(Y(1, s) = y(1) \mid Z) \prod_{t=2}^{\bar{T}} \Pr(Y(t, s) = y(t) \mid Y^{t-1}(s) = y^{t-1}, Z) \\ &= \Pr(Y(1, s) = y(1) \mid S(1) = s(1), Z) \prod_{t=2}^{\bar{T}} \Pr(Y(t, s) = y(t) \mid Y^{t-1}(s) = y^{t-1}, S^t = s^t, Z). \end{aligned}$$

Invoking (NA), in particular $Y(t, s) = Y_{p_0}(t)$ and $Y^{t-1}(s) = Y_{p_0}^{t-1}$ on $\{S^t = s^t\}$, produces

$$\begin{aligned} \Pr(Y(s) = y \mid Z) &= \Pr(Y_{p_0}(1) = y(1) \mid S(1) = s(1), Z) \\ &\quad \times \prod_{t=2}^{\bar{T}} \Pr(Y_{p_0}(t) = y(t) \mid Y_{p_0}^{t-1} = y^{t-1}, S^t = s^t, Z). \end{aligned} \quad (2)$$

This is a version of Robins' (1997) “ g -computation formula”.^{24,25} We can sequentially identify each component on the left hand side of the first expression, and hence identify the counterfactual

²⁴Gill and Robins (2001) present versions of (NA) and (M-1) for the case with more general distributions of treatments, and prove a version of the g -computation formula for the general case. For a random vector X and a function f that is integrable with respect to the distribution of X , let $\int_{x \in A} f(x) \Pr(X \in dx) = E[f(X)\mathbf{1}(X \in A)]$. Then,

$$\begin{aligned} \Pr(Y(s) \in A \mid Z) &= \int_{y \in A} \Pr(Y_{p_0}(\bar{T}) \in dy(\bar{T}) \mid Y_{p_0}^{\bar{T}-1} = y^{\bar{T}-1}, S^{\bar{T}} = s^{\bar{T}}, Z) \\ &\quad \vdots \\ &\quad \times \Pr(Y_{p_0}(2) \in dy(2) \mid Y_{p_0}(1) = y(1), S^2 = s^2, Z) \\ &\quad \times \Pr(Y_{p_0}(1) \in dy(1) \mid S(1) = s(1), Z), \end{aligned}$$

where A is a set of $Y(s)$. The right-hand side of this expression is almost surely unique under regularity conditions presented by Gill and Robins (2001).

²⁵An interesting special case arises if the outcomes are survival indicators, that is if $Y_{p_0}(t) = 1$ if the agent survives up to and including time t and $Y_{p_0}(t) = 0$ otherwise, $t \geq 1$. Then, no anticipation (NA) requires that treatment after death does not affect survival, and the g -computation formula simplifies considerably (Abbring, 2003).

distributions. This establishes identification of the distribution of $Y(s)$ by expressing it in terms of objects that can be identified from data. Identification is exact (or “tight”) in the sense that the identifying assumptions, no anticipation and sequential randomization, do not restrict the factual data and are therefore not testable (Gill and Robins, 2001).²⁶

Example 1. Consider a two-period ($\bar{T} = 2$) version of the model in which agents take either “treatment” (1) or “control” (0) in each period. Then, $S(1)$ and $S(2)$ have values in $\mathcal{S} = \{0, 1\}$. The potential outcomes in period t are $Y(t, (0, 0))$, $Y(t, (0, 1))$, $Y(t, (1, 0))$ and $Y(t, (1, 1))$. For example, $Y(2, (0, 0))$ is the outcome in period 2 in the case that the agent is assigned to the control group in each of the two periods. Using Bayes’ rule, it follows that

$$\Pr(Y(s) = y|Z) = \Pr(Y(1, s) = y(1) | Z) \Pr(Y(2, s) = y(2) | Y(1, s) = y(1), Z). \quad (3)$$

The g -computation approach to constructing $\Pr(Y(s) = y|Z)$ from data replaces the two probabilities in the right-hand side with probabilities of the observed (by the econometrician) variables (Y_{p_0}, S, Z) . First, note that $\Pr(Y(1, s) = y(1) | Z) = \Pr(Y(1, s) = y(1) | S(1) = s(1), Z)$ by (M-1). Moreover, (NA) ensures that potential outcomes in period 1 do not depend on the treatment status in period 2, so that

$$\Pr(Y(1, s) = y(1) | Z) = \Pr(Y_{p_0}(1) = y(1) | S(1) = s(1), Z).$$

Similarly, subsequently invoking (NA) and (M-1), then (M-1), and then (NA), gives

$$\begin{aligned} & \Pr(Y(2, s) = y(2) | Y(1, s) = y(1), Z) \\ &= \Pr(Y(2, s) = y(2) | Y_{p_0}(1), S(1) = s(1), Z) \quad (\text{by (NA) and (M-1)}) \\ &= \Pr(Y(2, s) = y(2) | Y_{p_0}(1), S = s, Z) \quad (\text{by (M-1)}) \\ &= \Pr(Y_{p_0}(2) = y(2) | Y_{p_0}(1), S = s, Z). \quad (\text{by (NA)}) \end{aligned}$$

Substituting these equations into the right-hand side of (3) gives the g -computation formula,

$$\Pr(Y(s) = y|Z) = \Pr(Y_{p_0}(1) = y(1) | S(1) = s(1), Z) \Pr(Y_{p_0}(2) = y(2) | Y_{p_0}(1) = y(1), S = s, Z).$$

Note that the right-hand side expression does not generally reduce to $\Pr(Y_{p_0} = y|S = s, Z)$. This would require the stronger, static matching condition $S \perp\!\!\!\perp Y(s) | Z$, which we have not assumed here.

Matching on pretreatment covariates is a special case of the g -computation approach. Suppose that the entire treatment path is assigned independently of potential outcomes given pretreatment covariates Z or, more precisely, $S \perp\!\!\!\perp Y(s) | Z$ for all s . This implies sequential randomization (M-1), and directly gives identification of the distributions of $Y(s)|Z$ and $Y(s)$. The matching assumption imposes no restriction on the data since $Y(s)$ is only observed if $S = s$. The no-anticipation condition

²⁶Gill and Robins’ (2001) analysis only involves causal inference on a final outcome (i.e., our $Y(s, \bar{T})$) and does not invoke the no-anticipation condition. However, their proof directly applies to the case studied in this chapter.

(NA) is not required for identification in this special case because no conditioning on S^t is required. Matching on pretreatment covariates is equivalent to matching in a static model. The distribution of $Y(s)|Z$ is identified without (NA), and assuming it to be true would impose testable restrictions on the data. In particular, it would imply that treatment assignment cannot be dependent on past outcomes given Z . The static matching assumption is not likely to hold in applications where treatment is dynamically assigned based on information on intermediate outcomes. This motivates an analysis based on the more subtle sequential randomization assumption. An alternative approach, developed in section 5, is to explicitly model and identify the evolution of the unobservables.

Gill and Robins claim that their sequential randomization and no-anticipation assumptions are “neutral”, “for free”, or “harmless”. As we will argue later, from an economic perspective, some of the model assumptions, notably the no-anticipation assumption, can be interpreted as substantial behavioral/informational assumptions. For example, Heckman and Vytlacil (2005, 2007b) and Heckman and Navarro (2004) show how matching imposes the condition that marginal and average returns are equal. Because of these strong assumptions, econometricians sometimes phrase their “neutrality” result more negatively as a non-identification result (Abbring and Van den Berg, 2003b), since it is possible that (M-1) and/or (NA) may not hold.

3.2 Policy Evaluation and Dynamic Discrete-Choice Analysis

3.2.1 The Effects of Policies

Consider a counterfactual policy p' such that the corresponding allocation of treatments $s_{p'}$ satisfies sequential randomization, as in (M-1):

(M-2) *For all treatment sequences s and all t ,*

$$s_{p'}(t) \perp\!\!\!\perp (Y(t, s), \dots, Y(\bar{T}, s)) \mid (Y_{p'}^{t-1}, s_{p'}^{t-1} = s^{t-1}, Z).$$

The treatment assignment rule $s_{p'}$ is equivalent to what Gill and Robins (2001) call a “randomized plan”. The outcome distribution under such a rule cannot be constructed by integrating the distributions of $\{Y(s)\}$ with respect to the distribution of $s_{p'}$, because there may be feedback from intermediate outcomes into treatment assignment. Instead, under the assumptions of the previous subsection and a support condition, we can use a version of the g -computation formula for randomized plans given by Gill and Robins to compute the distribution of outcomes under the policy

p' :²⁷

$$\begin{aligned} \Pr(Y_{p'} = y | Z) &= \sum_{s \in \mathcal{S}} \Pr(Y_{p_0}(1) = y(1) | S(1) = s(1), Z) \Pr(s_{p'}(1) = s(1) | Z) \\ &\quad \times \prod_{t=2}^{\bar{T}} [\Pr(Y_{p_0}(t) = y(t) | Y_{p_0}^{t-1} = y^{t-1}, S^t = s^t, Z) \\ &\quad \times \Pr(s_{p'}(t) = s(t) | Y_{p'}^{t-1} = y^{t-1}, s_{p'}^{t-1} = s^{t-1}, Z)] \end{aligned} \quad (4)$$

In the special case of static matching on Z , so that $s_{p'} \perp\!\!\!\perp U | Z$, this simplifies to integrating the distribution of $Y_{p_0} | (S = s, Z)$ over the distribution of $s_{p'} | Z$:²⁸

$$\Pr(Y_{p'} = y | Z) = \sum_{s \in \mathcal{S}} \Pr(Y_{p_0} = y | S = s, Z) \Pr(s_{p'} = s | Z).$$

3.2.2 Policy Choice and Optimal Policies

We now consider the problem of choosing a policy p that is optimal according to some criterion. This problem is both of normative interest and of descriptive interest if actual policies are chosen to be optimal. We could, for example, study the optimal assignment a' of constraints and incentives to agents. Alternatively, we could assume that agents pick τ to maximize their utilities, and use the methods discussed in this section to model τ .

Under the policy invariance assumptions that underlie the treatment-effect approach, p only affects outcomes through its implied treatment allocation s_p . Thus, the problem of choosing an optimal policy boils down to choosing an optimal treatment allocation s_p under informational and other constraints specific to the problem at hand. For example, suppose that the planner and the

²⁷The corresponding formula for the case with general treatment distributions is

$$\begin{aligned} \Pr(Y_{p'} \in A | Z) &= \int_{y \in A} \int_{s \in \mathcal{S}} \Pr(Y_{p_0}(\bar{T}) \in dy(\bar{T}) | Y_{p_0}^{\bar{T}-1} = y^{\bar{T}-1}, S^{\bar{T}} = s^{\bar{T}}, Z) \\ &\quad \times \Pr(s_{p'}(\bar{T}) \in ds(\bar{T}) | Y_{p'}^{\bar{T}-1} = y^{\bar{T}-1}, s_{p'}^{\bar{T}-1} = s^{\bar{T}-1}, Z) \\ &\quad \vdots \\ &\quad \times \Pr(Y_{p_0}(2) \in dy(2) | Y_{p_0}(1) = y(1), S(1) = s(1), Z) \\ &\quad \times \Pr(s_{p'}(2) \in ds(2) | Y_{p'}(1) = y(1), s_{p'}(1) = s(1), Z) \\ &\quad \times \Pr(Y_{p_0}(1) \in dy(1) | S(1) = s(1), Z) \Pr(s_{p'}(1) \in ds(1) | Z). \end{aligned}$$

The support condition on $s_{p'}$ requires that, for each t , the distribution of $s_{p'}(t) | (Y_{p'}^{t-1} = y^{t-1}, s_{p'}^{t-1} = s^{t-1}, Z = z)$ is absolutely continuous with respect to the distribution of $S(t) | (Y_{p_0}^{t-1} = y^{t-1}, S^{t-1} = s^{t-1}, Z = z)$ for almost all (y^{t-1}, s^{t-1}, z) from the distribution of $(Y_{p_0}^{t-1}, S^{t-1}, Z)$.

²⁸In the general case, this condition becomes

$$\Pr(Y_{p'} \in A | Z) = \int_{s \in \mathcal{S}} \Pr(Y_{p_0} \in A | S = s, Z) \Pr(s_{p'} \in ds | Z).$$

agents have the same information, $\mathcal{I}_P(p) = \mathcal{I}_A(p)$, the planner assigns eligibility to a program by a , and agents fully comply, so that $\mathcal{B} = \mathcal{S}$ and $s_p = a$. Then, s_p can be any rule from \mathcal{A} and is adapted to $\mathcal{I}_P(p) = \mathcal{I}_A(p)$.

For expositional convenience, we consider the optimal choice of a treatment assignment s_p adapted to the agent’s information $\mathcal{I}_A(p)$ constructed earlier. We will use the word “agents” to refer to the decision maker in this problem, even though it can also apply to the planner’s decision problem. An econometric approach to this problem is to estimate explicit dynamic choice models with explicit choice-outcome relationships. One emphasis in the literature is on Markovian discrete-choice models that satisfy Rust’s (1987) conditional-independence assumption (see Rust, 1994). Other assumptions are made in the literature and we exposit them in section 5.

Here, we explore the use of Rust’s (1987) model as a model of treatment choice in a dynamic treatment-effect setting. In particular, we make explicit the additional structure that Rust’s model, and in particular his conditional-independence assumption, imposes on Robins’ dynamic potential-outcomes model. We follow Rust (1987) and focus on a finite treatment (control) space \mathcal{S} . In the notation of our model, payoffs are determined by the outcomes Y_p , treatment choices s_p , the “cost shocks” V , and the covariates Z . Rust (1987) assumes that $\{Y_p(t-1), V_t, Z\}$ is a controlled first-order Markov process, with initial condition $Y_p(0) \equiv 0$ and control s_p .²⁹ As before, V_t and Z are not causally affected by choices, but $Y_p(t)$ may causally depend on current and past choices. The agents choose a treatment assignment rule s_p that maximizes

$$E \left[\sum_{t=1}^{\bar{T}} \Upsilon_t \{Y_p(t-1), V_t, s_p(t), Z\} + \Upsilon_{\bar{T}+1} \{Y_p(\bar{T}), Z\} \middle| \mathcal{I}_A(1) \right], \quad (5)$$

for some (net and discounted) utility functions Υ_t and $\mathcal{I}_A(1) = \mathcal{I}_A(1, p)$, which is independent of p . $\Upsilon_{\bar{T}+1} \{Y_p(\bar{T}), Z\}$ is the terminal value. Under standard regularity conditions on the utility functions, we can solve backward for the optimal policy s_p . Because of Rust’s Markov assumption, s_p has a Markovian structure,

$$s_p(t) \perp\!\!\!\perp (Y_p^{t-2}, V^{t-1}) \mid [Y_p(t-1), V_t, Z],$$

for $t = 2, \dots, \bar{T}$, and $\{Y_p(t-1), V_t, Z\}$ is a first-order Markov process. Note that Z enters the model as an observed (by the econometrician) factor that shifts net utility. A key assumption embodied in the specification of (5) is time-separability of utility. Rust (1987), in addition, imposes separability between observed and unobserved state variables. This assumption plays no essential role in expositing the core ideas in Rust, and we will not make it here.

Rust’s (1987) conditional-independence assumption imposes two key restrictions on the decision problem. It is instructive to consider these restrictions in isolation from Rust’s Markov restriction. We make the model’s causal structure explicit using the potential-outcomes notation. Note that the model has a recursive causal structure—the payoff-relevant state is controlled by current and past

²⁹Rust (1987) assumes an infinite-horizon, stationary environment. Here, we present a finite-horizon version to facilitate a comparison with the dynamic potential-outcomes model and to link up with the analysis in section 5.

choices only— and satisfies no-anticipation condition (NA). Setting $Y(0, s) \equiv 0$ for specificity, and ignoring the Markov restriction, Rust’s conditional-independence assumption requires, in addition to the assumption that there are no direct causal effects of choices on V , that

$$Y(s, t) \perp\!\!\!\perp V^t \mid [Y^{t-1}(s), Z], \text{ and} \quad (6)$$

$$V_{t+1} \perp\!\!\!\perp V^t \mid [Y^t(s), Z], \quad (7)$$

for all s and t . As noted by Rust (1987, p. 1011), condition (6) ensures that the observed (by the econometrician) controlled state evolves independently of the unobserved payoff-relevant variables. It is equivalent to (Florens and Mouchart, 1982)³⁰

$$(M-3) \quad [Y(s, t), \dots, Y(s, \bar{T})] \perp\!\!\!\perp V^t \mid [Y^{t-1}(s), Z] \text{ for all } t \text{ and } s.$$

In turn, (M-3) implies (M-1) and is equivalent to the assumption that (M-2) holds for all s_p .³¹

Condition (7) excludes serial dependence of the unobserved payoff-relevant variables conditional on past outcomes. In contrast, Robins’ g -computation framework allows for such serial dependence, provided that sequential randomization holds if serial dependence is present. For example, if $V \perp\!\!\!\perp U \mid Z$, then (M-1) and its variants hold without further assumptions on the time series structure of V_t .

The first-order Markov assumption imposes additional restrictions on potential outcomes. These restrictions are twofold. First, potential outcomes follow a first-order Markov process. Second, $s(t)$ only directly affects the Markov transition from $Y(t, s)$ to $Y(t + 1, s)$. This strengthens the no-anticipation assumption presented in section 3.1.1. The Markov assumption also requires that V_{t+1} only depends on $Y(s, t)$, and not on $Y^{t-1}(s)$, given $Y(s, t)$.

In applications, we may assume that actual treatment assignment S solves the Markovian decision problem. Together with specifications of Υ_t , this further restricts the dynamic choice-outcome model. Alternatively, one could make other assumptions on S and use (5) to define and find an optimal, and typically counterfactual, assignment rule s_p .

Our analysis shows that the substantial econometric literature on the structural empirical analysis of Markovian decision problems under conditional independence can be applied to policy evaluation under sequential randomization. Conversely, methods developed for potential-outcomes models with sequential randomization can be applied to learn about aspects of dynamic discrete-choice models. Murphy (2003) develops methods to estimate an optimal treatment assignment rule using Robins’ dynamic potential-outcomes model with sequential randomization (M-2).

3.3 The Information Structure of Policies

One concern about methods for policy evaluation based on the potential-outcomes model is that potential outcomes are sometimes reduced-form representations of dynamic models of agent’s choices.

³⁰Note that (6) is a Granger (1969) noncausality condition stating that, for all s and conditional on Z , V does not cause $Y(s)$.

³¹If V has redundant components, that is components that do not nontrivially enter any assignment rule s_p , (M-3) imposes more structure, but structure that is irrelevant to the decision problem and its empirical analysis.

A policy maker choosing optimal policies typically faces a population of agents who act on the available information, and their actions in turn affect potential outcomes. For example, in terms of the model of section 3.2.2, a policy may change financial incentives—the $b \in \mathcal{B}$ assigned through a could enter the net utilities Υ_t —and leave it to the agents to control outcomes by choosing treatment. In econometric policy evaluation, it is therefore important to carefully model the information \mathcal{I}_A that accumulates to the agents in different program states and under different policies, separately from the policy maker’s information \mathcal{I}_P .

This can be contrasted with common practice in biostatistics. Statistical analyses of the effects of drugs on health are usually concerned with the physician’s (planner’s) information and decision problem. Gill and Robins’ (2001) sequential randomization assumption, for example, is often justified by the assumption that physicians base their treatment decisions on observable (by the analyst) information only. This literature, however, often ignores the possibility that many variables known to the physician may not be known to the observing statistician and that the agents being given drugs alter the protocols.

Potential outcomes will often depend on the agent’s information. Failure to correctly model the information will often lead to violation of (NA) and failure of invariance. Potential outcomes may therefore not be valid inputs in a policy evaluation study. A naive specification of potential outcomes would only index treatments by actual participation in, e.g., job search assistance or training programs. Such a naive specification is incomplete in the context of economies inhabited by forward-looking agents who make choices that affect outcomes. In specifying potential outcomes, we should not only consider the effects of actual program participation, but also the effects of the information available to agents about the program and policy. We now illustrate this point.

Example 2. Black, Smith, Berger, and Noel (2003) analyze the effect of compulsory training and employment services provided to unemployment insurance (UI) claimants in Kentucky on the exit rate from UI and earnings. In the program they study, letters are sent out to notify agents some time ahead whether they are selected to participate in the program. This information is recorded in a database and available to them. They can analyze the letter as part of a program that consists of information provision and subsequent participation in training. The main empirical finding of their paper is that the threat of future mandatory training conveyed by the letters is more effective in increasing the UI exit rate than training itself.

The data used by Black, Smith, Berger, and Noel (2003) are atypical of many economic data sets, because the data collectors carefully record the information provided to agents. This allows Black et al. to analyze the effects of the provision of information along with the effects of actual program participation. In many econometric applications, the information on the program under study is less rich. Data sets may provide information on actual participation in training programs and some background information on how the program is administered. Typically, however, the data do not record all of the letters sent to agents and do not record every phone conversation between administrators and agents. Then, the econometrician needs to make assumptions on how this information accumulates for agents. In many applications, knowledge of specific institutional mechanisms of assignment can be used to justify specific informational assumptions.

Example 3. Abbring, Van den Berg, and Van Ours (2005) analyze the effect of punitive benefits reductions, or sanctions, in Dutch UI on re-employment rates. In the Netherlands, UI claimants have to comply with certain rules concerning search behavior and registration. If a claimant violates these rules, a sanction may be applied. A sanction is a punitive reduction in benefits for some period of time and may be accompanied by increased levels of monitoring by the UI agency.³² Abbring, Van den Berg, and Van Ours (2005) use administrative data and know the re-employment duration, the duration at which a sanction is imposed if a sanction is imposed, and some background characteristics for each UI case.

Without prior knowledge of the Dutch UI system, an analyst might make a variety of informational assumptions. One extreme is that UI claimants know at the start of their UI spells that their benefits will be reduced at some specific duration if they are still claiming UI at that duration. This results in a UI system with entitlement periods that are tailored to individual claimants and that are set and revealed at the start of the UI spells. In this case, claimants will change their labor market behavior from the start of their UI spell in response to the future benefits reduction (e.g., Mortensen, 1977). At another extreme, claimants receive no prior signals of impending sanctions and there are no anticipatory effects of actual benefits reductions. However, agents may still be aware of the properties of the sanctions process and to some extent this will affect their behavior. Abbring, Van den Berg, and Van Ours (2005) analyze a search model with these features. Abbring and Van den Berg (2003b) provide a structural example where the data cannot distinguish between these two informational assumptions. Abbring, Van den Berg, and Van Ours (2005) use institutional background information to argue in favor of the second informational assumption as the one that characterizes their data.

If data on information provision are not available and simplifying assumptions on the program’s information structure cannot be justified, the analyst needs to model the information that accumulates to agents as an unobserved determinant of outcomes. This is the approach followed, and further discussed, in section 5.

The information determining outcomes typically includes aspects of the policy. In Example 2, the letter announcing future training will be interpreted differently in different policy environments. If agents are forward looking, the letter will be more informative under a policy that specifies a strong relation between the letter and mandatory training in the population than under a policy that allocates letters and training independently. In Example 3, the policy is a monitoring regime. Potential outcomes are UI durations under different sanction times. A change in monitoring policy changes the value of unemployment. In a job-search model with forward-looking agents, agents will respond by changing their search effort and reservation wages, and UI duration outcomes will change. In either example, potential outcomes are not invariant to variation in the policy. In the terminology of Hurwicz (1962), the policy is not “structural” with regard to potential outcomes and violates the invariance assumptions presented in Heckman and Vytlačil (2007a). One must control for the effects of agents’ information.

³²See Grubb (2000) for a review of sanction systems in the OECD.

3.4 Selection on Unobservables

In econometric program evaluations, (sequentially) randomized assignment is unlikely to hold. We illustrate this in the models developed in section 5. Observational data are characterized by a lot of heterogeneity among agents, as documented by the empirical examples in Abbring and Heckman (2007) and in Heckman, LaLonde, and Smith (1999). This heterogeneity is unlikely to be fully captured by the observed variables in most data sets. In a dynamic context, such unmeasured heterogeneity leads to violations of the assumptions of Gill and Robins (2001) and Rust (1987) that choices represent a sequential randomization. This is true even if the unmeasured variables only affect the availability of slots in programs but not outcomes directly. If agents are rational, forward-looking and observe at least some of the unmeasured variables that the econometrician does not, they will typically respond to these variables through their choice of treatment and investment behavior. In this case, the sequential randomization condition fails.

For the same reason, identification based on instrumental variables is relatively hard to justify in dynamic models (Hansen and Sargent, 1980; Rosenzweig and Wolpin, 2000; Abbring and Van den Berg, 2005). If the candidate instruments only vary across persons but not over time for the same person, then they are not likely to be valid instruments because they affect expectations and future choices and may affect current potential outcomes. Instead of using instrumental variables that vary only across persons, we require instruments based on unanticipated person-specific shocks that affect treatment choices but not outcomes at each point in time. In the context of continuously assigned treatments, the implied data requirements seem onerous. To achieve identification, Abbring and Van den Berg (2003b) focus on regressor variation rather than exclusion restrictions in a sufficiently smooth model of continuous-time treatment effects. We discuss their analysis in section 4. Heckman and Navarro (2007) show that curvature conditions, not exclusion restrictions, that result in the same variables having different effects on choices and outcomes in different periods, are motivated by economic theory and can be exploited to identify dynamic treatment effects in discrete time without literally excluding any variables. We discuss their analysis in section 5. We now consider a formulation of the analysis in continuous time.

4 The Event-History Approach to Policy Analysis

The discrete-time models just discussed in section 3 have an obvious limitation. Time is continuous and many events are best described by a continuous-time model. There is a rich field of continuous-time event-history analysis that has been adapted to conduct policy evaluation analysis.³³ For example, the effects of training and counseling on unemployment durations and job stability have been analyzed by applying event-history methods to data on individual labor-market and training histories (Ridder, 1986; Card and Sullivan, 1988; Gritz, 1993; Ham and LaLonde, 1996; Eberwein et al., 1997; Bonnal et al., 1997). Similarly, the moral hazard effects of unemployment insurance have been studied by analyzing the effects of time-varying benefits on labor-market transitions (e.g.,

³³Abbring and Van den Berg (2004) discuss the relation between the event-history approach to program evaluation and more standard latent-variable and panel-data methods, with a focus on identification issues.

Meyer, 1990; Abbring et al., 2005; Van den Berg et al., 2004). In fields like epidemiology, the use of event-history models to analyze treatment effects is widespread (see, e.g., Andersen et al., 1993; Keiding, 1999).

The event-history approach to program evaluation is firmly rooted in the econometric literature on state dependence (lagged dependent variables) and heterogeneity (Heckman and Borjas, 1980; and Heckman, 1981a). Event-history models along the lines of Heckman and Singer (1984, 1986) are used to jointly model transitions into programs and transitions into outcome states. Causal effects of programs are modelled as the dependence of individual transition rates on the individual history of program participation. Dynamic selection effects are modelled by allowing for dependent unobserved heterogeneity in both the program and outcome transition rates.

Without restrictions on the class of models considered, true state dependence and dynamic selection effects cannot be distinguished.³⁴ Any history dependence of current transition rates can be explained both as true state dependence and as the result of unobserved heterogeneity that simultaneously affects the history and current transitions. This is a dynamic manifestation of the problem of drawing causal inference from observational data. In applied work, researchers avoid this problem by imposing additional structure. A typical, simple, example is a mixed semi-Markov model in which the causal effects are restricted to program participation in the previous spell (e.g., Bonnal, Fougère, and Sérandon, 1997). There is a substantial literature on the identifiability of state-dependence effects and heterogeneity in duration and event-history models that exploit such additional structure (see Heckman and Taber, 1994, and Van den Berg, 2001 for reviews). Here, we provide discussion of some canonical cases.

4.1 Treatment Effects in Duration Models

4.1.1 Dynamically Assigned Binary Treatments and Duration Outcomes

We first consider the simplest case of mutual dependence of events in continuous time, involving only two binary events. This case is sufficiently rich to capture the effect of a dynamically assigned binary treatment on a duration outcome. Binary events in continuous time can be fully characterized by the time at which they occur and a structural model for their joint determination is a simultaneous-equations model for durations. We develop such a model along the lines of Abbring and Van den Berg (2003b). This model is an extension, with general marginal distributions and general causal and spurious dependence of the durations, of Freund's (1961) bivariate exponential model.

Consider two continuously-distributed random durations Y and S . We refer to one of the durations, S , as the time to treatment and to the other duration, Y , as the outcome duration. Such an asymmetry arises naturally in many applications. For example, in Abbring, Van den Berg, and Van Ours's (2005) study of unemployment insurance, the treatment is a punitive benefits reduction (sanction) and the outcome re-employment. The re-employment process continues after imposition of a sanction, but the sanctions process is terminated by re-employment. The current exposition,

³⁴See Heckman and Singer (1986).

however, is symmetric and unifies both cases. It applies to both the asymmetric setup of the sanctions example and to applications in which both events may causally affect the other event.

Let $Y(s)$ be the potential outcome duration that would prevail if the treatment time is externally set to s . Similarly, let $S(y)$ be the potential treatment time resulting from setting the outcome duration to y . We assume that *ex ante* heterogeneity across agents is fully captured by observed covariates X and unobserved covariates V , assumed to be external and temporally invariant. Treatment causally affects the outcome duration through its hazard rate. We denote the hazard rate of $Y(s)$ at time t for an agent with characteristics (X, V) by $\theta_Y(t|s, X, V)$. Similarly, outcomes affect the treatment times through its hazard $\theta_S(t|y, X, V)$. Causal effects on hazard rates are produced by recursive economic models driven by point processes, such as search models. We provide an example below, and further discussion in section 4.3.

Without loss of generality, we partition V into (V_S, V_Y) and assume that $\theta_Y(t|s, X, V) = \theta_Y(t|s, X, V_Y)$ and $\theta_S(t|y, X, V) = \theta_S(t|y, X, V_S)$. Intuitively, V_S and V_Y are the unobservables affecting, respectively, treatment and outcome, and the joint distribution of (V_S, V_Y) is unrestricted. In particular, V_S and V_Y may have elements in common.

The corresponding integrated hazard rates are defined by $\Theta_Y(t | s, X, V_Y) = \int_0^t \theta_Y(u | s, X, V_Y) du$ and $\Theta_S(t | y, X, V_S) = \int_0^t \theta_S(u | y, X, V_S) du$. For expositional convenience, we assume that these integrated hazards are strictly increasing in t . We also assume that they diverge to ∞ as $t \rightarrow \infty$, so that the duration distributions are non-defective.³⁵ Then, $\Theta_Y(Y(s) | s, X, V_Y)$ and $\Theta_S(S(y) | y, X, V_S)$ are unit exponential for all $y, s \in \mathbb{R}_+$.³⁶ This implies the following model of potential outcomes and treatments,³⁷

$$Y(s) = y(s, X, V_Y, \varepsilon_Y) \text{ and } S(y) = s(y, X, V_S, \varepsilon_S),$$

for some unit exponential random variables ε_Y and ε_S that are independent of (X, V) , $y = \Theta_Y^{-1}$, and $s = \Theta_S^{-1}$.

The exponential errors ε_Y and ε_S embody the *ex post* shocks that are inherent to the individual hazard processes, that is the randomness in the transition process after conditioning on covariates X and V and survival. We assume that $\varepsilon_Y \perp\!\!\!\perp \varepsilon_S$, so that $\{Y(s)\}$ and $\{S(y)\}$ are only dependent through the observed and unobserved covariates (X, V) . This conditional-independence assumption is weaker than the conditional-independence assumption underlying the analysis of section 3 and

³⁵ Abbring and Van den Berg (2003b) allow for defective distributions, which often have structural interpretations. For example, some women never have children and some workers will never leave a job. See Abbring (2002) for discussion.

³⁶ Let $T | X$ be distributed with density $f(t|X)$, non-defective cumulative distribution function $F(t|X)$, and hazard rate $\theta(t|X) = f(t|X)/[1 - F(t|X)]$. Then, $\int_0^T \theta(t|X) dt = -\ln[1 - F(T|X)]$ is a unit exponential random variable that is independent of X .

³⁷ The causal hazard model only implies that the distributions of ε_Y and ε_S are invariant across assigned treatments and outcomes, respectively; their realizations may not be. This is sufficient for the variation of $y(s, X, V_Y, \varepsilon_Y)$ with s and of $s(y, X, V_S, \varepsilon_S)$ with y to have a causal interpretation. The further restriction that the random variables ε_Y and ε_S are invariant is made for simplicity, and is empirically innocuous. See Abbring and Van den Berg (2003b) for details and Freedman (2004) for discussion.

used in matching, because it allows for conditioning on the invariant unobservables V . It shares this feature with the discrete-time models developed in section 5.

We assume a version of the no-anticipation condition of section 3.1.1: for all $t \in \mathbb{R}_+$,

$$\theta_Y(t|s, X, V_Y) = \theta_Y(t|s', X, V_Y) \text{ and } \theta_S(t|y, X, V_S) = \theta_S(t|y', X, V_S)$$

for all $s, s', y, y' \in [t, \infty)$. This excludes effects of anticipation of the treatment on the outcome. Similarly, there can be no anticipation effects of future outcomes on the treatment time hazard.

Example 4. Consider a standard search model describing the job search behavior of an unemployed individual (e.g., Mortensen, 1986) with characteristics (X, V) . Job offers arrive at a rate $\lambda > 0$ and are random draws from a given distribution F . Both λ and F may depend on (X, V) , but, for notational simplicity, we suppress all explicit representations of conditioning on (X, V) throughout this example. An offer is either accepted or rejected. A rejected offer cannot be recalled at a later time. The individual initially receives a constant flow of unemployment-insurance benefits. However, the individual faces the risk of a sanction—a permanent reduction of his benefits to some lower, constant level—at some point during his unemployment spell. During the unemployment spell, sanctions arrive independently of the job-offer process at a constant rate $\mu > 0$. The individual cannot foresee the exact time a sanction is imposed, but he knows the distribution of these times.³⁸ The individual chooses a job-acceptance rule so as to maximize his expected discounted lifetime income. Under standard conditions, this is a reservation-wage rule: at time t , the individual accepts each wage of $\underline{w}(t)$ or higher. The corresponding re-employment hazard rate is $\lambda(1 - F(\underline{w}(t)))$. Apart from the sanction, which is not foreseen and arrives at a constant rate during the unemployment spell, the model is stationary. This implies that the reservation wage is constant, say equal to \underline{w}_0 , up to and including time s , jumps to some lower level $\underline{w}_1 < \underline{w}_0$ at time s and stays constant at \underline{w}_1 for the remainder of the unemployment spell if benefits would be reduced at time s .

The model is a version of the simultaneous-equations model for durations. To see this, let Y be the re-employment duration and S the sanction time. The potential-outcome hazards are

$$\theta_Y(t|s) = \begin{cases} \lambda_0 & \text{if } 0 \leq t \leq s \\ \lambda_1 & \text{if } t > s, \end{cases}$$

where $\lambda_0 = \lambda[1 - F(\underline{w}_0)]$ and $\lambda_1 = \lambda[1 - F(\underline{w}_1)]$, and clearly $\lambda_1 \geq \lambda_0$. Similarly, the potential-treatment time hazards are $\theta_S(t|y) = \mu$ if $0 \leq t \leq y$, and 0 otherwise. Note that the no-anticipation condition follows naturally from the recursive structure of the economic decision problem in this case, in which we have properly accounted for all relevant components of agent information sets. Furthermore, the assumed independence of the job offer and sanction processes at the individual level for given (X, V) implies that $\varepsilon_Y \perp\!\!\!\perp \varepsilon_S$.

The actual outcome and treatment are related to the potential outcomes and treatments by $S = S(Y)$ and $Y = Y(S)$. The no-anticipation assumption ensures that this system has a unique solution

³⁸This is a rudimentary version of the search model with punitive benefits reductions, or sanctions, of Abbring, Van den Berg, and Van Ours (2005). The main difference is that in the present version of the model the sanctions process cannot be controlled by the agent.

(Y, S) by imposing a recursive structure on the underlying transition processes. Without anticipation effects, current treatment and outcome hazards only depend on past outcome and treatment events, and the transition processes evolve recursively (Abbring and Van den Berg, 2003b). Together with a distribution $G(\cdot | X)$ of $V | X$, this gives a non-parametric structural model of the distribution of $(Y, S) | X$ that embodies general simultaneous causal dependence of Y and S , dependence of (Y, X) on observed covariates X , and general dependence of the unobserved errors V_Y and V_S .

There are two reasons for imposing further restrictions on this model. First, it is not identified from data on (Y, S, X) . Take a version of the model with selection on unobservables ($V_Y \not\perp V_S | X$) and consider the distribution of $(Y, S)|X$ generated by this version of the model. Then, there exists an alternative version of the model that satisfies both no-anticipation and $V_Y \perp V_S | X$, and that generates the same distribution of $(Y, S)|X$ (Abbring and Van den Berg, 2003b, Proposition 1). In other words, for each version of the model with selection on unobservables and anticipation effects, there is an observationally-equivalent model version that satisfies no-anticipation and conditional randomization. This is a version of the nonidentification result discussed in section 3.1.

Second, even if we ensure nonparametric identification by assuming no-anticipation and conditional randomization, we cannot learn about the agent-level causal effects embodied in y and s without imposing even further restrictions. At best, under regularity conditions we can identify $\theta_Y(t|s, X) = E[\theta_Y(t|s, X, V_Y)|X, Y(s) \geq t]$ and $\theta_S(t|y, X) = E[\theta_S(t|y, X, V_S)|X, S(y) \geq t]$ from standard hazard regressions (e.g., Andersen et al., 1993; Fleming and Harrington, 1991). Thus, we can identify the distributions of $Y(s)|X$ and $S(y)|X$, and therefore solve the selection problem if we are only interested in these distributions. However, if we are also interested in the causal effects on the corresponding hazard rates for given X, V , we face an additional dynamic selection problem. The hazards of the identified distributions of $Y(s)|X$ and $S(y)|X$ only condition on observed covariates X , and not on unobserved covariates V , and are confounded with dynamic selection effects (Heckman and Borjas, 1980; Heckman and Singer, 1986; Meyer, 1996; Abbring and Van den Berg, 2005). For example, the difference between $\theta_Y(t|s, X)$ and $\theta_Y(t|s', X)$ does not only reflect agent-level differences between $\theta_Y(t|s, X, V_Y)$ and $\theta_Y(t|s', X, V_Y)$, but also differences in the subpopulations of survivors $\{X, Y(s) \geq t\}$ and $\{X, Y(s') \geq t\}$ on which the hazards are computed.

In the next two subsections, we discuss what can be learned about treatment effects in duration models under additional model restrictions. We take the no-anticipation assumption as fundamental. As explained in section 3, this requires that we measure and include in our model all relevant information needed to define potential outcomes. However, we relax the randomization assumption. We first consider Abbring and Van den Berg's (2003b) analysis of identification without exclusion restrictions. They argue that these results are useful, because exclusion restrictions are hard to justify in an inherently dynamic setting with forward-looking agents. Abbring and Van den Berg (2005) further clarify this issue by studying inference for treatment effects in duration models using a social experiment. We discuss what can be learned from such experiments at the end of this section.

4.1.2 Identifiability Without Exclusion Restrictions

Abbring and Van den Berg consider an extension of the multivariate Mixed Proportional Hazard (MPH) model (Lancaster, 1979) in which the hazard rates of $Y(s) \mid (X, V)$ and $S(y) \mid (X, V)$ are given by

$$\theta_Y(t \mid s, X, V) = \begin{cases} \lambda_Y(t)\phi_Y(X)V_Y & \text{if } t \leq s \\ \lambda_Y(t)\phi_Y(X)\delta_Y(t, s, X)V_Y & \text{if } t > s \end{cases} \quad (8)$$

and

$$\theta_S(t \mid y, X, V) = \begin{cases} \lambda_S(t)\phi_S(X)V_S & \text{if } t \leq y \\ \lambda_S(t)\phi_S(X)\delta_S(t, y, X)V_S & \text{if } t > y, \end{cases} \quad (9)$$

respectively, and $V = (V_S, V_Y)$ is distributed independently of X . The baseline hazards $\lambda_Y : \mathbb{R}_+ \rightarrow (0, \infty)$ and $\lambda_S : \mathbb{R}_+ \rightarrow (0, \infty)$ capture duration dependence of the individual transition rates. The integrated hazards are $\Lambda_Y(t) = \int_0^t \lambda_Y(\tau) d\tau < \infty$ and $\Lambda_S(t) = \int_0^t \lambda_S(\tau) d\tau < \infty$ for all $t \in \mathbb{R}_+$. The regressor functions $\phi_Y : \mathcal{X} \rightarrow (0, \infty)$ and $\phi_S : \mathcal{X} \rightarrow (0, \infty)$ are assumed to be continuous, with $\mathcal{X} \subset \mathbb{R}^q$ the support of X . In empirical work, these functions are frequently specified as $\phi_Y(x) = \exp(x'\beta_Y)$ and $\phi_S(x) = \exp(x'\beta_S)$ for some parameter vectors β_Y and β_S . We will not make such parametric assumptions. Note that the fact that both regressor functions are defined on the same domain \mathcal{X} is not restrictive, because each function ϕ_Y and ϕ_S can “select” certain elements of X by being trivial functions of the other elements. In the parametric example, the vector β_Y would only have nonzero elements for those regressors that matter to the outcome hazard. The functions δ_Y and δ_S capture the causal effects. Note that $\delta_Y(t, s, X)$ only enters $\theta_Y(t \mid s, X, V)$ at durations $t > s$, so that the model satisfies no anticipation of treatment assumption (NA). Similarly, it satisfies no anticipation of outcomes and has a recursive causal structure as required by the no-anticipation assumption. If $\delta_Y = 1$, treatment is ineffective; if δ_Y is larger than 1, it stochastically reduces the remaining outcome duration.

Note that this model allows δ_Y and δ_S to depend on elapsed duration t , past endogenous events, and the observed covariates X , but not on V . Abbring and Van den Berg also consider an alternative model that allows δ_Y and δ_S to depend on unobservables in a general way, but not on past endogenous events.

Abbring and Van den Berg show that these models are nonparametrically identified from single-spell data under the conditions for the identification of competing-risks models based on the multivariate MPH model given by Abbring and Van den Berg (2003a). Among other conditions are the requirements that there is some independent local variation of the regressor effects in both hazard rates and a finite-mean restriction on V , which are standard in the analysis of multivariate MPH models. With multiple-spell data, most of these assumptions, and the MPH structure, can be relaxed (Abbring and Van den Berg, 2003b).

The models can be parameterized in a flexible way and estimated by maximum likelihood. Typical parameterizations involve linear-index structures for the regressor and causal effects, a discrete distribution G , and piecewise-constant baseline hazards λ_S and λ_Y . Abbring and Van den Berg (2003c) develop a simple graphical method for inference on the sign of $\ln(\delta_Y)$ in the absence

of regressors. Abbring, Van den Berg, and Van Ours (2005) present an empirical application.

4.1.3 Inference Based on Instrumental Variables

The concerns expressed in subsection 3.4 about the validity of exclusion restrictions in dynamic settings carry over to event-history models.

Example 5. A good illustration of this point is offered by the analysis of Eberwein, Ham, and LaLonde (1997), who study the effects of a training program on labor-market transitions. Their data are particularly nice, as potential participants are randomized into treatment and control groups at some baseline point in time. This allows them to estimate the effect of intention to treat (with training) on subsequent labor-market transitions. This is directly relevant to policy evaluation in the case that the policy involves changing training enrollment through offers of treatment which may or may not be accepted by agents.

However, Eberwein et al. are also interested in the effect of actual participation in the training program on post program labor-market transitions. This is a distinct problem, because compliance with the intention-to-treat protocol is imperfect. Some agents in the control group are able to enroll in substitute programs, and some agents in the treatment group choose never to enroll in a program at all. Moreover, actual enrollment does not take place at the baseline time, but is dispersed over time. Those in the treatment group are more likely to enroll earlier. This fact, coupled with the initial randomization, suggests that the intention-to-treat indicator might be used as an instrument for identifying the effect of program participation on employment and unemployment spells.

The dynamic nature of enrollment into the training program, and the event-history focus of the analysis complicate matters considerably. Standard instrumental-variables methods cannot be directly applied. Instead, Eberwein et al. use a parametric duration model for pre and post program outcomes that excludes the intention-to-treat indicator from directly determining outcomes. They specify a duration model for training enrollment that includes an intention-to-treat indicator as an explanatory variable, and specify a model for labor-market transitions that excludes the intention-to-treat indicator and imposes a no-anticipation condition on the effect of actual training participation on labor-market transitions. Such a model is consistent with an environment in which agents cannot perfectly foresee the actual training time they will be assigned and in which they do not respond to information about this time revealed by their assignment to an intention-to-treat group. This is a strong assumption. In a search model with forward-looking agents, for example, such information would typically affect the *ex ante* values of unemployment and employment. Then, it would affect the labor-market transitions before actual training enrollment through changes in search efforts and reservation wages, unless these are both assumed to be exogenous. An assumption of perfect foresight on the part of the agents being studied only complicates matters further.

Abbring and Van den Berg (2005) study what can be learned about dynamically assigned programs from social experiments if the intention-to-treat instrument cannot be excluded from the outcome equation. They develop bounds, tests for unobserved heterogeneity, and point-identification

results that extend those discussed in this section.³⁹

4.2 Treatment Effects in More General Event-History Models

It is instructive to place the causal duration models developed in section 4.1 in the more general setting of event-history models with state dependence and heterogeneity. We do this following Abbring's (2008) analysis of the mixed semi-Markov model.

4.2.1 The Mixed Semi-Markov Event-History Model

The model is formulated in a fashion that is analogous to the frameworks of Heckman and Singer (1986). The point of departure is a continuous-time stochastic process assuming values in a finite set \mathcal{S} at each point in time. We will interpret realizations of this process as agents' event histories of transitions between states in the state space \mathcal{S} .

Suppose that event histories start at real-valued random times T_0 in an \mathcal{S} -valued random state S_0 , and that subsequent transitions occur at random times T_1, T_2, \dots such that $T_0 < T_1 < T_2 < \dots$. Let S_l be the random destination state of the transition at T_l . Taking the sample paths of the event-history process to be right-continuous, we have that S_l is the state occupied in the interval $[T_l, T_{l+1})$.

Suppose that heterogeneity among agents is captured by vectors of time-constant observed covariates X and unobserved covariates V .⁴⁰ In this case, state dependence in the event-history process for given individual characteristics X, V has a causal interpretation.⁴¹ We structure such state dependence by assuming that the event-history process conditional on X, V is a time-homogeneous semi-Markov process. Conditional on X, V the length of a spell in a state and the destination state of the transition ending that spell depend only on the past through the current state. In our notation, $(\Delta T_l, S_l) \perp\!\!\!\perp \{(T_i, S_i), i = 0, \dots, l-1\} \mid S_{l-1}, X, V$, where $\Delta T_l = T_l - T_{l-1}$ is the length of spell l . Also, the distribution of $(\Delta T_l, S_l) \mid S_{l-1}, X, V$ does not depend on l . Note that, conditional on X, V , $\{S_l, l \geq 0\}$ is a time-homogeneous Markov chain under these assumptions.

Non-trivial dynamic selection effects arise because V is not observed. The event-history process conditional on observed covariates X only is a mixed semi-Markov process. If V affects the initial state S_0 , or transitions from it, subpopulations of agents in different states at some time t typically have different distributions of the unobserved characteristics V . Therefore, a comparison of the subsequent transitions in two such subpopulations does not only reflect state dependence, but also sorting of agents with different unobserved characteristics into the different states they occupy at time t .

³⁹In the special case that a static treatment, or treatment plan, is assigned at the start of the spell, standard instrumental-variables methods may be applied. See Abbring and Van den Berg (2005).

⁴⁰We restrict attention to time-invariant observed covariates for expositional convenience. The analysis can easily be adapted to more general time-varying external covariates. Restricting attention to time-constant regressors is a worst-case scenario for identification. External time variation in observed covariates aids identification (Heckman and Taber, 1994).

⁴¹We could make this explicit by extending the potential-outcomes model of section 4.1.2 to the general event-history setup. However, this would add a lot of complexity, but little extra insight.

We model $\{(\Delta T_l, S_l), l \geq 1\} | T_0, S_0, X, V$ as a repeated competing-risks model. Due to the mixed semi-Markov assumption, the latent durations corresponding to transitions into the possible destination states in the l th spell only depend on the past through the current state S_{l-1} , conditional on X, V . This implies that we can fully specify the repeated competing-risks model by specifying a set of origin-destination-specific latent durations, with corresponding transition rates. Let T_{jk}^l denote the latent duration corresponding to the transition from state j to state k in spell l . We explicitly allow for the possibility that transitions between certain (ordered) pairs of states may be impossible. To this end, define the correspondence $\mathcal{Q} : \mathcal{S} \rightarrow \sigma(\mathcal{S})$ assigning to each $s \in \mathcal{S}$ the set of all destination states to which transitions are made from s with positive probability.⁴² Here, $\sigma(\mathcal{S})$ is the set of all subsets of \mathcal{S} (the “power set” of \mathcal{S}). Then, the length of spell l is given by $\Delta T_l = \min_{s \in \mathcal{Q}(S_{l-1})} T_{S_{l-1}s}^l$, and the destination state by $S_l = \arg \min_{s \in \mathcal{Q}(S_{l-1})} T_{S_{l-1}s}^l$.

We take the latent durations to be mutually independent, jointly independent of T_0, S_0 , and identically distributed across spells l , all conditional on X, V . This reflects both the mixed semi-Markov assumption and the additional assumption that all dependence between the latent durations corresponding to the competing risks in a given spell l is captured by the observed regressors X and the unobservables V . This is a standard assumption in econometric duration analysis, which, with the semi-Markov assumption, allows us to characterize the distribution of $\{(\Delta T_l, S_l), l \geq 1\} | T_0, S_0, X, V$ by specifying origin-destination-specific hazards $\theta_{jk}(t | X, V)$ for the marginal distributions of $T_{jk}^l | X, V$.

We assume that the hazards $\theta_{jk}(t | X, V)$ are of the mixed proportional hazard (MPH) type:⁴³

$$\theta_{jk}(t | X, V) = \begin{cases} \lambda_{jk}(t) \phi_{jk}(X) V_{jk} & \text{if } k \in \mathcal{Q}(j) \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

The baseline hazards $\lambda_{jk} : \mathbb{R}_+ \rightarrow (0, \infty)$ have integrated hazards $\Lambda_{jk}(t) = \int_0^t \lambda_{jk}(\tau) d\tau < \infty$, for all $t \in \mathbb{R}_+$. The regressor functions $\phi_{jk} : \mathcal{X} \rightarrow (0, \infty)$ are assumed to be continuous. Finally, the $(0, \infty)$ -valued random variable V_{jk} is the scalar component of V that affects the transition from state j to state k . Note that we allow for general dependence between the components of V . This way, we can capture, for example, that agents with lower re-employment rates have higher training enrolment rates.

This model fully characterizes the distribution of the transitions $\{(\Delta T_l, S_l), l \geq 1\}$ conditional on the initial conditions T_0, S_0 and the agent’s characteristics X, V . A complete model of the event histories $\{(T_l, S_l), l \geq 0\}$ conditional on X, V would in addition require a specification of the initial conditions T_0, S_0 for given X, V . It is important to stress here that T_0, S_0 are the initial conditions of the event-history process itself, and should not be confused with the initial conditions in a particular sample (which we will discuss below). In empirical work, interest in the dependence between start times T_0 and characteristics X, V is often limited to the observation that the distribution of agents’

⁴²Throughout this section, we assume that \mathcal{Q} is known. It is important to note, however, that \mathcal{Q} can actually be identified trivially in all cases considered.

⁴³Proportionality can be relaxed if we have data on sufficiently long event-histories. See Honoré (1993) and Abbring and Van den Berg (2003a,b) for related arguments for various multi-spell duration models.

characteristics may vary over cohorts indexed by T_0 . The choice of initial state S_0 may in general be of some interest, but is often trivial. For example, we could model labor-market histories from the calendar time T_0 at which agents turn 15 onwards. In an economy with perfect compliance to a mandatory schooling up to age 15, the initial state S_0 would be “(mandatory) schooling” for all. Therefore, we will not consider a model of the event history’s initial conditions, but instead focus on the conditional model of subsequent transition histories.

Because of the semi-Markov assumption, the distribution of $\{(\Delta T_l, S_l), l \geq 1\} | T_0, S_0, X, V$ only depends on S_0 , and not T_0 . Thus, T_0 only affects observed event histories through cohort effects on the distribution of unobserved characteristics V . The initial state S_0 , on the other hand, may both have causal effects on subsequent transitions and be informative on the distribution of V . For expositional clarity, we assume that $V \perp\!\!\!\perp (T_0, S_0, X)$. This is true, for example, if all agents start in the same state, so that S_0 is degenerate, and V is independent of the start date T_0 and the observed covariates X .

An econometric model for transition histories conditional on the observed covariates X can be derived from the model of $\{(\Delta T_l, S_l), l \geq 1\} | S_0, X, V$ by integrating out V . The exact way this should be done depends on the sampling scheme used. Here, we focus on sampling from the population of event-histories. We assume that we observe the covariates X , the initial state S_0 , and the first \bar{L} transitions from there. Then, we can model these transitions for given S_0, X by integrating the conditional model over the distribution of V .

Abbring (2008) discusses more complex, and arguably more realistic, sampling schemes. For example, when studying labor-market histories we may randomly sample from the stock of the unemployed at a particular point in time. Because the unobserved component V affects the probability of being unemployed at the sampling date, the distribution of $V | X$ in the stock sample does not equal its population distribution. This is again a dynamic version of the selection problem. Moreover, in this case, we typically do not observe an agent’s entire labor-market history from T_0 onwards. Instead, we may have data on the time spent in unemployment at the sampling date and on labor-market transitions for some period after the sampling date. This “initial-conditions problem” complicates matters further (Heckman, 1981b).

In the next two subsections, we first discuss some examples of applications of the model and then review a basic identification result for the simple sampling scheme above.

4.2.2 Applications to Program Evaluation

Several empirical papers study the effect of a single treatment on some outcome duration or set of transitions. Two approaches can be distinguished. In the first approach, the outcome and treatment processes are explicitly and separately specified. The second approach distinguishes treatment as one state within a single event-history model with state dependence.

The first approach is used in a variety of papers in labor economics. Eberwein, Ham, and LaLonde (1997) specify a model for labor market transitions in which the transition intensities between various labor market states (not including treatment) depend on whether someone has been assigned to a training program in the past or not. Abbring, Van den Berg, and Van Ours (2005)

and Van den Berg, Van der Klaauw, and Van Ours (2004) specify a model for re-employment durations in which the re-employment hazard depends on whether a punitive benefits reduction has been imposed in the past. Similarly, Van den Berg, Holm, and Van Ours (2002) analyze the duration up to transition into medical trainee positions and the effect of an intermediate transition into a medical assistant position (a “stepping-stone job”) on this duration. In all of these papers, the outcome model is complemented with a hazard model for treatment choice.

These models fit into the framework of section 4.1.2 or a multi-state extension thereof. We can rephrase the class of models discussed in section 4.1.2 in terms of a simple event-history model with state-dependence as follows. Distinguish three states, untreated (O), treated (P) and the exit state of interest (E), so that $\mathcal{S} = \{O, P, E\}$. All subjects start in O , so that $S_0 = O$. Obviously, we do not want to allow for all possible transitions between these three states. Instead, we restrict the correspondence \mathcal{Q} representing the possible transitions as follows:

$$\mathcal{Q}(s) = \begin{cases} \{P, E\} & s = O, \\ \{E\} & \text{if } s = P, \\ \emptyset & s = E. \end{cases}$$

State dependence of the transition rates into E captures treatment effects in the sense of section 4.1.2. Not all models in Abbring and Van den Berg (2003b) are included in the semi-Markov setup discussed here. In particular, in this paper we do not allow the transition rate from P to E to depend on the duration spent in O . This extension with “lagged duration dependence” (Heckman and Borjas, 1980) would be required to capture one variant of their model.

The model for transitions from “untreated” (O) is a competing risks model, with program enrolment (transition to P) and employment (E) competing to end the untreated spell. If the unobservable factor V_{OE} that determines transitions to employment and the unobservable factor V_{OP} affecting program enrolment are dependent, then program enrolment is selective in the sense that the initial distribution of V_{OE} — and also typically that of V_{PE} — among those who enroll at a given point in time does not equal its distribution among survivors in O up to that time.⁴⁴

The second approach is used by Gritz (1993) and Bonnal, Fougère, and Sérandon (1997), among others. Consider the following simplified setup. Suppose workers are either employed (E), unemployed (O), or engaged in a training program (P). We can now specify a transition process among these three labor market states in which a causal effect of training on unemployment and employment durations is modeled as dependence of the various transition rates on the past occurrence of a training program in the labor market history. Bonnal, Fougère, and Sérandon (1997) only have limited information on agents’ labor-market histories before the sample period. Partly to avoid difficult initial-conditions problems, they restrict attention to “first order lagged occurrence dependence” (Heckman and Borjas, 1980) by assuming that transition rates only depend on the current and previous states occupied. Such a model is not directly covered by the semi-Markov model, but with a simple augmentation of the state space it can be covered. In particular, we

⁴⁴Note that, in addition, the survivors in O themselves are a selected subpopulation. Because V affects survival in O , the distribution of V among survivors in O is not equal to its population distribution.

have to include lagged states in the state space on which the transition process is defined. Because there is no lagged state in the event-history's first spell, initial states should be defined separately. So, instead of just distinguishing states in $\mathcal{S}^* = \{E, O, P\}$, we distinguish augmented states in $\mathcal{S} = \{(s, s') \in (\mathcal{S}^* \cup \{I\}) \times \mathcal{S}^* : s \neq s'\}$. Then, (I, s) , $s \in \mathcal{S}^*$, denote the initial states, and $(s, s') \in \mathcal{S}$ the augmented state of an agent who is currently in s' and came from $s \neq s'$. In order to preserve the interpretation of the model as a model of lagged occurrence dependence, we have to exclude certain transitions by specifying

$$\mathcal{Q}(s, s') = \{(s', s''), s'' \in \mathcal{S}^* \setminus \{s'\}\}.$$

This excludes transitions to augmented states that are labeled with a lagged state different from the origin state. Also, it ensures that agents never return to an initial state. For example, from the augmented state (O, P) —previously unemployed and currently enrolled in a program—only transitions to augmented states (P, s'') —previously enrolled in a program and currently in s'' —are possible. Moreover, it is not possible to be currently employed and transiting to initially unemployed, (I, O) . Rather, an employed person who loses her job would transit to (E, O) —currently unemployed and previously employed.

The effects are now modeled as simple state-dependence effects. For example, the effect of training on the transition rate from unemployment to employment is simply the contrast between the individual transition rate from (E, O) to (O, E) and the transition rate from (P, O) to (O, E) . Dynamic selection into the augmented states (E, O) and (P, O) , as specified by the transition model, confounds the empirical analysis of these training effects. Note that due to the fact that we have restricted attention to first-order lagged occurrence dependence, there are no longer-run effects of training on transition rates from unemployment to employment.

4.2.3 Identification Without Exclusion Restrictions

In this section, we state a basic identification result for the following sampling scheme. Suppose that the economist randomly samples from the population of event-histories, and that we observe the first \bar{L} transitions (including destinations) for each sampled event-history, with the possibility that $\bar{L} = \infty$.⁴⁵ Thus, we observe a random sample of $\{(T_l, S_l), l \in \{0, 1, \dots, \bar{L}\}\}$, and X .

First note that we can only identify the determinants of θ_{jk} for transitions (j, k) that occur with positive probability among the first \bar{L} transitions. Moreover, without further restrictions, we can only identify the joint distribution of a vector of unobservables corresponding to (part of) a sequence of transitions that can be observed among the first \bar{L} transitions.

With this qualification, identification can be proved by extending Abbring and Van den Berg's (2003a) analysis of the MPH competing risks model to the present setting. This analysis assumes that transition rates have an MPH functional form. Identification again requires specific moments of V to be finite, and independent local variation in the regressor effects.

⁴⁵Note that this assumes away econometric initial-conditions problems of the type previously discussed.

4.3 A Structural Perspective

Without further restrictions, the causal duration model of section 4.1.1 is versatile. It can be generated as the reduced form of a wide variety of continuous-time economic models driven by point processes. Leading examples are sequential job-search models in which job-offer arrival rates, and other model parameters, depend on agent characteristics (X, V) and policy interventions (see, e.g., Mortensen, 1986, and Example 4).

The MPH restriction on this model, however, is hard to justify from economic theory. In particular, nonstationary job-search models often imply interactions between duration and covariate effects; the MPH model only results under strong assumptions (Heckman and Singer, 1986; Van den Berg, 2001). Similarly, an MPH structure is hard to generate from models in which agents learn about their individual value of the model’s structural parameters, that is about (X, V) , through Bayesian updating.

An alternative class of continuous-time models, not discussed in this chapter, specifies durations as the first time some Gaussian or more general process crosses a threshold. Such models are closely related to a variety of dynamic economic models. They have attracted recent attention in statistics (see, e.g., Aalen and Gjessing, 2004). Abbring (2007) analyzes identifiability of “mixed hitting-time models”, continuous-time threshold-crossing models in which the parameters depend on observed and unobserved covariates, and discusses their link with optimizing models in economics. This is a relatively new area of research, and a full development is beyond the scope of this paper. It extends to a continuous-time framework the dynamic threshold crossing model developed in Heckman (1981a,b) that is used in the next subsection of this chapter.

We now discuss a complementary discrete-time approach where it is possible to make many important economic distinctions that are difficult to make in the setting of continuous-time models and to avoid some difficult measure-theoretic problems. In the structural version, it is possible to specify precisely agent information sets in a fashion that is not possible in conventional duration models.

5 Dynamic Discrete Choice and Dynamic Treatment Effects

Heckman and Navarro (2007) and Cunha, Heckman, and Navarro (2007) present econometric models for analyzing time to treatment and the consequences of the choice of a particular treatment time. Treatment may be a medical intervention, stopping schooling, opening a store, conducting an advertising campaign at a given date or renewing a patent. Associated with each treatment time, there can be multiple outcomes. They can include a vector of health status indicators and biomarkers; lifetime employment and earnings consequences of stopping at a particular grade of schooling; the sales revenue and profit generated from opening a store at a certain time; the revenues generated and market penetration gained from an advertising campaign; or the value of exercising an option at a given time. Heckman and Navarro (2007) unite and contribute to the literatures on dynamic discrete choice and dynamic treatment effects. For both classes of models, they present semiparametric identification analyses. We summarize their work in this section. It is formulated in discrete

time, which facilitates the specification of richer unobserved and observed covariate processes than those entertained in the continuous-time framework of Abbring and Van den Berg (2003b).

Heckman and Navarro extend the literature on treatment effects to model choices of treatment times and the consequences of choice and link the literature on treatment effects to the literature on precisely formulated structural dynamic discrete-choice models generated from index models crossing thresholds. They show the value of precisely formulated economic models in extracting the information sets of agents, in providing model identification, in generating the standard treatment effects and in enforcing the nonanticipating behavior condition (NA) discussed in section 3.1.⁴⁶

They establish the semiparametric identifiability of a class of dynamic discrete-choice models for stopping times and associated outcomes in which agents sequentially update the information on which they act. They also establish identifiability of a new class of reduced-form duration models that generalize conventional discrete-time duration models to produce frameworks with much richer time series properties for unobservables and general time-varying observables and patterns of duration dependence than conventional duration models. Their analysis of identification of these generalized models requires richer variation driven by observables than is needed in the analysis of the more restrictive conventional models. However, it does not require conventional period-by-period exclusion restrictions, which are often difficult to justify. Instead, they rely on curvature restrictions across the index functions generating the durations that can be motivated by dynamic economic theory.⁴⁷ Their methods can be applied to a variety of outcome measures including durations.

The key to their ability to identify structural models is that they supplement information on stopping times or time to treatment with additional information on measured consequences of choices of time to treatment as well as measurements. The dynamic discrete-choice literature surveyed in Rust (1994) and Magnac and Thesmar (2002) focuses on discrete-choice processes with general preferences and state vector evolution equations, typically Markovian in nature. Rust's 1994 paper contains negative results on nonparametric identification of discrete-choice processes. Magnac and Thesmar (2002) present some positive results on nonparametric identification if certain parameters or distributions of unobservables are assumed to be known. Heckman and Navarro (2007) produce positive results on nonparametric identification of a class of dynamic discrete-choice models based on expected income maximization developed in labor economics by Flinn and Heckman (1982), Keane and Wolpin (1997) and Eckstein and Wolpin (1999). These frameworks are dynamic versions of the Roy model. Heckman and Navarro (2007) show how use of cross-equation restrictions joined with data on supplementary measurement systems can undo Rust's nonidentification result. We exposit their work and the related literature in this section. With their structural framework, they can distinguish objective outcomes from subjective outcomes (valuations by the decision maker) in a dynamic setting. Applying their analysis to health economics, they can identify the causal effects on health of a medical treatment as well as the associated subjective pain and suffering of a treatment

⁴⁶Aakvik, Heckman, and Vytlacil (2005), Heckman, Tobias, and Vytlacil (2001, 2003), Carneiro, Hansen, and Heckman (2001, 2003) and Heckman and Vytlacil (2005) show how standard treatment effects can be generated from structural models.

⁴⁷See Heckman and Honoré (1989) for examples of such an identification strategy in duration models. See also Cameron and Heckman (1998).

regime for the patient.⁴⁸ Attrition decisions also convey information about agent preferences about treatment.⁴⁹

They do not rely on the assumption of conditional independence of unobservables with outcomes, given observables, that is used throughout much of the dynamic discrete-choice literature and the dynamic treatment literature surveyed in section 3.⁵⁰ As noted in section 2, sequential conditional independence assumptions underlie recent work on reduced-form dynamic treatment effects.⁵¹ The semiparametric analysis of Heckman and Navarro (2007) based on factors generalizes matching to a dynamic setting. In their paper, some of the variables that would produce conditional independence and would justify matching if they were observed, are treated as unobserved match variables. They are integrated out and their distributions are identified.⁵² They consider two classes of models. We review both.

5.1 Semiparametric Duration Models and Counterfactuals

Heckman and Navarro (2007), henceforth HN, develop a semiparametric index model for dynamic discrete choices that extends conventional discrete-time duration analysis. They separate out duration dependence from heterogeneity in a semiparametric framework more general than conventional discrete-time duration models. They produce a new class of reduced-form models for dynamic treatment effects by adjoining time-to-treatment outcomes to the duration model. This analysis builds on Heckman (1981a,b,c).

Their models are based on a latent variable for choice at time s ,

$$I(s) = \Psi(s, Z(s)) - \eta(s),$$

where the $Z(s)$ are observables and $\eta(s)$ are unobservables from the point of view of the econometrician. Treatments at different times may have different outcome consequences which they model after analyzing the time to treatment equation. Define $D(s)$ as an indicator of receipt of treatment at date s . Treatment is taken the first time $I(s)$ becomes positive. Thus,

$$D(s) = \mathbf{1}[I(s) \geq 0, I(s-1) < 0, \dots, I(1) < 0],$$

where the indicator function $\mathbf{1}[\cdot]$ takes the value of 1 if the term inside the braces is true.⁵³ They derive conditions for identifying a model with general forms of duration dependence in the time to treatment equation using a large sample from the distribution of (D, Z) .

⁴⁸See Chan and Hamilton (2006) for a structural dynamic empirical analysis of this problem.

⁴⁹See Heckman and Smith (1998). Use of participation data to infer preferences about outcomes is developed in Heckman (1974).

⁵⁰See, e.g., Rust (1987), Manski (1993), Hotz and Miller (1993) and the papers cited in Rust (1994).

⁵¹See, e.g., Gill and Robins (2001) and Lechner and Miquel (2002).

⁵²For estimates based on this idea, see Carneiro, Hansen, and Heckman (2003), Aakvik, Heckman, and Vytlačil (2005), Cunha and Heckman (2007, 2008); Cunha, Heckman, and Navarro (2005, 2006), and Heckman and Navarro (2005).

⁵³This framework captures the essential feature of any stopping time model. For example, in a search model with one wage offer per period, $I(s)$ is the gap between market wages and reservation wages at time s . See, e.g., Flinn and Heckman (1982). This framework can also approximate the explicit dynamic discrete-choice model analyzed in section 5.2.

5.1.1 Single-Spell Duration Model

Individuals are assumed to start spells in a given (exogenously determined) state and to exit the state at the beginning of time period S .⁵⁴ S is thus a random variable representing total completed spell length. Let $D(s) = 1$ if the individual exits at time s , $S = s$, and $D(s) = 0$ otherwise. In an analysis of drug treatments, S is the discrete-time period in the course of an illness at the beginning of which the drug is administered. Let $\bar{S} (< \infty)$ be the upper limit on the time the agent being studied can be at risk for a treatment. It is possible in this example that $D(1) = 0, \dots, D(\bar{S}) = 0$, so that a patient never receives treatment. In a schooling example, “treatment” is not schooling, but rather dropping out of schooling.⁵⁵ In this case, \bar{S} is an upper limit to the number of years of schooling, and $D(\bar{S}) = 1$ if $D(1) = 0, \dots, D(\bar{S} - 1) = 0$.

The duration model can be specified recursively in terms of the threshold-crossing behavior of the sequence of underlying latent indices $I(s)$. Recall that $I(s) = \Psi(s, Z(s)) - \eta(s)$, with $Z(s)$ being the regressors that are observed by the analyst. The $Z(s)$ can include expectations of future outcomes given current information in the case of models with forward-looking behavior. For a given stopping time s , let $D^s = (D(1), \dots, D(s))$ and designate by $d(s)$ and d^s values that $D(s)$ and D^s assume. Thus, $d(s)$ can be zero or one and d^s is a sequence of s zeros or a sequence containing $s - 1$ zeros and a single one. Denote a sequence of all zeros by (0) , regardless of its length. Then,

$$\begin{aligned} D(1) &= \mathbf{1}[I(1) \geq 0] \\ \text{and} & \\ D(s) &= \begin{cases} \mathbf{1}[I(s) \geq 0] & \text{if } D^{s-1} = (0) \\ 0 & \text{otherwise,} \end{cases} \quad s = 2, \dots, \bar{S}. \end{aligned} \tag{11}$$

For $s = 2, \dots, \bar{S}$, the indicator $\mathbf{1}[I(s) \geq 0]$ is observed if and only if the agent is still at risk of treatment, $D^{s-1} = (0)$. To identify period s parameters from period s outcomes, one must condition on all past outcomes and control for any selection effects.

Let $Z = (Z(1), \dots, Z(\bar{S}))$, and let $\eta = (\eta(1), \dots, \eta(\bar{S}))$. Assume that Z is statistically independent of η . Heckman and Navarro (2007) assume that $\Psi(s, Z(s)) = Z(s)\gamma_s$. We deal with a more general case. $\Psi(Z) = (\Psi(1, Z(1)), \dots, \Psi(\bar{S}, Z(\bar{S})))$. We let Ψ denote the abstract parameter. Depending on the values assumed by $\Psi(s, Z(s))$, one can generate very general forms of duration dependence that depend on the values assumed by the $Z(s)$. HN allow for period-specific effects of regressors on the latent indices generating choices.

This model is the reduced form of a general dynamic discrete-choice model. Like many reduced-form models, the link to choice theory is not clearly specified. It is not a conventional multinomial choice model in a static (perfect certainty) setting with associated outcomes.

⁵⁴Thus we abstract from the initial-conditions problem discussed in Heckman (1981b).

⁵⁵In the drug treatment example, S may designate the time a treatment regime is completed.

5.1.2 Identification of Duration Models with General Error Structures and Duration Dependence

Heckman and Navarro (2007) establish semiparametric identification of the model of equation (11) assuming access to a large sample of i.i.d. (D, Z) observations. Let $Z^s = (Z(1), \dots, Z(s))$. Data on (D, Z) directly identify the conditional probability $\Pr(D(s) = d(s) | Z^s, D^{s-1} = (0))$ *a.e.* $F_{Z^s | D^{s-1}=(0)}$ where $F_{Z^s | D^{s-1}=(0)}$ is the distribution of Z^s conditional on previous choices $D^{s-1} = (0)$. Assume that $(\Psi, F_\eta) \in \Phi \times \mathcal{H}$, where F_η is the distribution of η and $\Phi \times \mathcal{H}$ is the parameter space. The goal is to establish conditions under which knowledge of $\Pr(D(s) = d(s) | Z, D^{s-1} = (0))$ *a.e.* $F_{Z | D^{s-1}=(0)}$ allows the analyst to identify a unique element of $\Phi \times \mathcal{H}$. They use a limit strategy that allows them to recover the parameters by conditioning on large values of the indices of the preceding choices. This identification strategy is widely used in the analysis of discrete choice.⁵⁶

They establish sufficient conditions for the identification of model (11). We prove the following more general result:

Theorem 1. *For the model defined by equation (11), assume the following conditions:*

- (i) $\eta \perp\!\!\!\perp Z$.
- (ii) η is an absolutely continuous random variable on $\mathbb{R}^{\bar{S}}$ with support $\prod_{s=1}^{\bar{S}} (\underline{\eta}(s), \bar{\eta}(s))$, where $-\infty \leq \underline{\eta}(s) < \bar{\eta}(s) \leq +\infty$, for all $s = 1, \dots, \bar{S}$.
- (iii) The $\Psi(s, Z(s))$ satisfy the Matzkin (1992) conditions for identification of nonparametric binary choice models, $s = 1, \dots, \bar{S}$.⁵⁷
- (iv) $\text{Supp}(\Psi^{s-1}(Z), Z(s)) = \text{Supp}(\Psi^{s-1}(Z)) \times \text{Supp}(Z(s))$, $s = 2, \dots, \bar{S}$.
- (v) $\text{Supp}(\Psi(Z)) \supseteq \text{Supp}(\eta)$.

Then F_η and $\Psi(Z)$ are identified, where the $\Psi(s, Z(s))$, $s = 1, \dots, \bar{S}$, are identified over the relevant support admitted by (ii).

Proof. We sketch the proof for $\bar{S} = 2$. The result for general \bar{S} follows by a recursive application of this argument. Consider the following three probabilities.

- (a) $\Pr(D(1) = 1 | Z = z) = \int_{\underline{\eta}(1)}^{\Psi(1, z(1))} f_{\eta(1)}(u) du$
- (b) $\Pr(D(2) = 1, D(1) = 0 | Z = z) = \int_{\underline{\eta}(2)}^{\Psi(2, z(2))} \int_{\Psi(1, z(1))}^{\bar{\eta}(1)} f_{\eta(1), \eta(2)}(u_1, u_2) du_1 du_2.$
- (c) $\Pr(D(2) = 0, D(1) = 0 | Z = z) = \int_{\Psi(2, z(2))}^{\bar{\eta}(2)} \int_{\Psi(1, z(1))}^{\bar{\eta}(1)} f_{\eta(1), \eta(2)}(u_1, u_2) du_1 du_2.$

⁵⁶See, e.g., Manski (1988), Heckman (1990), Heckman and Honoré (1989, 1990), Matzkin (1992, 1993), Taber (2000), and Carneiro, Hansen, and Heckman (2003). A version of the strategy of this proof was first used in psychology where agent choice sets are eliminated by experimenter manipulation. The limit set argument effectively uses regressors to reduce the choice set confronting agents. See Falmagne (1985) for a discussion of models of choice in psychology.

⁵⁷See Abbring and Heckman (2007, Appendix B.1) for a review of the conditions Matzkin (1992) imposes for identification of nonparametric binary choice models. See also Matzkin (1994).

The left-hand sides are observed from data on those who stop in period 1 (a); those who stop in period 2 (b); and those who terminate in the “0” state in period 2 (c). From Matzkin (1992), we can identify $\Psi(1, z(1))$ and $F_{\eta(1)}$ from (a). Using (b), we can fix $z(2)$ and vary $\Psi(1, z(1))$. From (iv) and (v), there exists a limit set \tilde{Z}_1 , possibly dependent on $z(2)$, such that $\lim_{z(1) \rightarrow \tilde{Z}_1} \Psi(1, z(1)) = \underline{\eta}(1)$. Thus we can construct

$$\Pr(D(2) = 0 \mid Z = z) = \int_{\Psi(2, z(2))}^{\bar{\eta}(2)} f_{\eta(2)}(u_2) du_2$$

and identify $\Psi(2, z(2))$ and $F_{\eta(2)}$. Using the $\Psi(1, z(1))$, $\Psi(2, z(2))$, one can trace out the joint distribution $F_{\eta(1), \eta(2)}$ over its support. Under the Matzkin conditions, identification is achieved on a non-negligible set. The proof generalizes in a straightforward way to general \tilde{S} . \square

Observe that if the $\eta(s)$ are bounded by finite upper and lower limits, we can only determine the $\Psi(s, Z(s))$ over the limits so defined. Consider the first step of the proof. Under the Matzkin conditions, $F_{\eta(1)}$ is known. From assumption (ii), we can determine

$$\Psi(1, z(1)) = F_{\eta(1)}^{-1}(\Pr(D(1) = 1 \mid Z = z)),$$

but only over the support $(\underline{\eta}(1), \bar{\eta}(1))$. If the support of $\eta(1)$ is \mathbb{R} , we determine $\Psi(1, z(1))$ for all $z(1)$. Heckman and Navarro (2007) analyze the special case $\Psi(s, Z(s)) = Z(s)\gamma_s$ and invoke sequential rank conditions to identify γ_s , even over limited supports. They also establish that the limit sets are non-negligible in this case so that standard definitions of identifiability (see, e.g., Matzkin, 1992) will be satisfied.⁵⁸ Construction of the limit set $\tilde{Z}_s, s = 1, \dots, \bar{S}$, depends on the functional form specified for the $\Psi(s, Z(s))$. For the linear-in-parameters case $\Psi(s, Z(s)) = Z(s)\gamma_s$, they are obtained by letting arguments get big or small. Matzkin (1992) shows how to establish the limit sets for functions in her family of functions.

A version of Theorem 1 with $\Psi(s, Z(s)) = Z(s)\gamma_s$ that allows dependence between Z and η^s except for one component can be proved using the analysis of Lewbel (2000) and Honoré and Lewbel (2002).⁵⁹

The assumptions of Theorem 1 will be satisfied if there are transition-specific exclusion restrictions for Z with the required properties. As noted in section 4, in models with many periods, this may be a demanding requirement. Very often, the Z variables are time invariant and so cannot be used as exclusion restrictions. Corollary 1 in HN, for the special case $\Psi(s, Z(s)) = Z(s)\gamma_s$, tells us that the HN version of the model can be identified, even if there are no conventional exclusion restrictions and the $Z(s)$ are the *same* across all time periods, if sufficient structure is placed on how the γ_s vary with s . Variations in the values of γ_s across time periods arise naturally in finite horizon

⁵⁸Heckman and Navarro (2007) prove their theorem for a model where $D(s) = \mathbf{1}[I(s) \leq 0]$ if $D^{s-1} = (0), s = 2, \dots, \bar{S}$. Our formulation of their result is consistent with the notation in this chapter.

⁵⁹HN discuss a version of such an extension at their website. Lewbel’s conditions are very strong. To account for general forms of dependence between Z and η^s requires modeling the exact form of the dependence. Nonparametric solutions to this problem remain an open question in the literature on dynamic discrete choice. One solution is to assume functional forms for the error terms, but in general, this is not enough to identify the model without further restrictions imposed. See Heckman and Honoré (1990).

dynamic discrete-choice models where a shrinking horizon produces different effects of the same variable in different periods. For example, in Wolpin’s (1987) analysis of a search model, the value function depends on time and the derived decision rules weight the same invariant characteristics differently in different periods. In a schooling model, parental background and resources may affect education continuation decisions differently at different stages of the schooling decision. The model generating equation (11) can be semiparametrically identified without transition-specific exclusions if the duration dependence is sufficiently general. For a proof, see Corollary 1 in Heckman and Navarro (2007).

The conditions of Theorem 1 are somewhat similar to the conditions on the regressor effects needed for identification of the continuous-time event-history models in section 4. One difference is that the present analysis requires independent variation of the regressor effects over the support of the distribution of the unobservables generating outcomes. The continuous-time analysis based on the functional form of the mixed proportional hazard model (MPH) as analyzed by Abbring and Van den Berg (2003a) only requires local independent variation.

Theorem 1 and Corollary 1 in HN have important consequences. The $\Psi(s, Z(s))$, $s = 1, \dots, \bar{S}$, can be interpreted as duration dependence parameters that are modified by the $Z(s)$ and that vary across the spell in a more general way than is permitted in mixed proportional hazards (MPH), generalized accelerated failure time (GAFT) models or standard discrete-time hazard models.⁶⁰ Duration dependence in conventional specifications of duration models is usually generated by variation in model intercepts. The regressors are allowed to interact with the duration dependence parameters. In the specifications justified by Theorem 1, the “heterogeneity” distribution F_η is identified for a general model. No special “permanent-transitory” structure is required for the unobservables although that specification is traditional in duration analysis. Their explicit treatment of the stochastic structure of the duration model is what allows HN to link in a general way the unobservables generating the duration model to the unobservables generating the outcome equations that are introduced in the next section. Such an explicit link is not currently available in the literature on continuous-time duration models for treatment effects surveyed in section 4, and is useful for modelling selection effects in outcomes across different treatment times. Their outcomes can be both discrete and continuous and are not restricted to be durations.

Under the conditions given in Corollary 1 of HN, no period-specific exclusion conditions are required on the Z . Hansen and Sargent (1980) and Abbring and Van den Berg (2003b) note that period-specific exclusions are not natural in reduced-form duration models designed to approximate forward-looking life cycle models. Agents make current decisions in light of their forecasts of future constraints and opportunities, and if they forecast some components well, and they affect current decisions, then they are in $Z(s)$ in period s . Corollary 1 in HN establishes identification without such exclusions. HN adjoin a system of counterfactual outcomes to their model of time to treatment to produce a model for dynamic counterfactuals. We summarize that work next.

⁶⁰See Ridder (1990) for a discussion of these models.

5.1.3 Reduced-Form Dynamic Treatment Effects

This section reviews a reduced-form approach to generating dynamic counterfactuals developed by HN. They apply and extend the analysis of Carneiro, Hansen, and Heckman (2003) and Cunha, Heckman, and Navarro (2005, 2006) to generate *ex post* potential outcomes and their relationship with the time to treatment indices $I(s)$ analyzed in the preceding subsection. With reduced-form models, it is difficult to impose restrictions from economic theory or to make distinctions between *ex ante* and *ex post* outcomes. In the structural model developed below, these and other distinctions can be made easily.

Associated with each treatment time s , $s = 1, \dots, \bar{S}$, is a vector of \bar{T} outcomes,

$$Y(s, X, U(s)) = (Y(1, s, X, U(1, s)), \dots, Y(t, s, X, U(t, s)), \dots, Y(\bar{T}, s, X, U(\bar{T}, s))) .$$

Outcomes depend on covariates X and $U(s) = (U(1, s), \dots, U(t, s), \dots, U(\bar{T}, s))$ that are, respectively, observable and unobservable by the econometrician. Elements of $Y(s, X, U(s))$ are outcomes associated with stopping or receiving treatment at the beginning of period s . They are factual outcomes if treatment s is actually selected ($S = s$ and $D(s) = 1$). Outcomes corresponding to treatments s' that are not selected ($D(s') = 0$) are counterfactuals. The outcomes associated with each treatment may be different, and indeed the treatments administered at different times may be different.

The components $Y(t, s, X, U(t, s))$ of the vector $Y(s, X, U(s))$ can be interpreted as the outcomes revealed at age t , $t = 1, \dots, \bar{T}$, and may themselves be vectors. The reduced-form approach presented in this section is not sufficiently rich to capture the notion that agents revise their anticipations of components of $Y(s, X, U(s))$, $s = 1, \dots, \bar{S}$, as they acquire information over time. This notion is systematically developed using the structural model discussed below in section 5.2.

The treatment “times” may be stages that are not necessarily connected with real times. Thus s may be a schooling level. The correspondence between stages and times is exact if each stage takes one period to complete. Our notation is more flexible, and time and periods can be defined more generally. Our notation in this section accommodates both cases.

Henceforth, whenever we have random variables with multiple arguments $R_0(s, Q_0, \dots)$ or $R_1(t, s, Q_0, \dots)$ where the argument list begins with treatment state s or both age t and state s (perhaps followed by other arguments Q_0, \dots), we will make use of several condensed notations: (a) dropping the first argument as we collect the components into vectors $R_0(Q_0, \dots)$ or $R_1(s, Q_0, \dots)$ of length \bar{S} or \bar{T} , respectively, and (b) going further in the case of R_1 , dropping the s argument as we collect the vectors $R_1(s, Q_0, \dots)$ into a single $\bar{S} \times \bar{T}$ array $R_1(Q_0, \dots)$, but also (c) suppressing one or more of the other arguments and writing $R_1(t, s)$ or $R_1(t, s, Q_0)$ instead of $R_1(t, s, Q_0, Q_1, \dots)$, etc. This notation is sufficiently rich to represent the life cycle of outcomes for persons who receive treatment at s . Thus, in a schooling example, the components of this vector may include life cycle earnings, employment, and the like associated with a person with characteristics $X, U(s)$, $s = 1, \dots, \bar{S}$, who completes s years of schooling and then forever ceases schooling. It could include earnings while in treatment at some level for persons who will eventually attain further schooling as well as post-school earnings.

We measure age and treatment time on the same time scale, with origin 1, and let $\bar{T} \geq \bar{S}$. Then, the $Y(t, s, X, U(t, s))$ for $t < s$ are outcomes realized while the person is in school at age t (s is the time the person will leave school; t is the current age) and before “treatment” (stopping schooling) has occurred. When $t \geq s$, these are post-school outcomes for treatment with s years of schooling. In this case, $t - s$ is years of post-school experience. In the case of a drug trial, the $Y(t, s, X, U(t, s))$ for $t < s$ are measurements observed before the drug is taken at s and if $t \geq s$, they are the post-treatment measurements.

Following Carneiro, Hansen, and Heckman (2003), the variables in $Y(t, s, X, U(t, s))$ may include discrete, continuous or mixed discrete-continuous components. For the discrete or mixed discrete-continuous cases, HN assume that latent continuous variables cross thresholds to generate the discrete components. Durations can be generated by latent index models associated with each outcome crossing thresholds analogous to the model presented in equation (11). In this framework, for example, we can model the effect of attaining s years of schooling on durations of unemployment or durations of employment.

The reduced-form analysis in this section does not impose restrictions on the temporal (age) structure of outcomes across treatment times in constructing outcomes and specifying identifying assumptions. Each treatment time can have its own age path of outcomes pre and post treatment. Outcomes prior to treatment and outcomes after treatment are treated symmetrically and both may be different for different treatment times. In particular, HN can allow earnings at age t for people who receive treatment at some future time s' to differ from earnings at age t for people who receive treatment at some future time s'' , $\min(s', s'') > t$ even after controlling for U and X .

This generality is in contrast with the analyses of Robins (1997) and Gill and Robins (2001) discussed in section 3 and the analysis of Abbring and Van den Berg (2003b) discussed in section 4. These analyses require exclusion of such anticipation effects to secure identification, because their models attribute dependence of treatment on past outcomes to selection effects. The sequential randomization assumption (M-1) underlying the work of Gill and Robins allows treatment decisions $S(t)$ at time t to depend on past outcomes $Y_{p_0}^{t-1}$ in a general way. Therefore, without additional restrictions, it is not possible to also identify *causal* (anticipatory) effects of treatment $S(t)$ on $Y_{p_0}^{t-1}$. The no-anticipation condition (NA) excludes such effects and secures identification in their framework.⁶¹ It is essential for applying the conditional independence assumptions in deriving the

⁶¹The role of the no-anticipation assumption in Abbring and Van den Berg (2003b) is similar. However, their main analysis assumes an asymmetric treatment-outcome setup in which treatment is not observed if it takes place after the outcome transition. In that case, the treatment time is censored at the outcome time. In this asymmetric setup, anticipatory effects of treatment on outcomes cannot be identified because the econometrician cannot observe variation of outcome transitions with future treatment times. This point may appear to be unrelated to the present discussion, but it is not. As was pointed out by Abbring and Van den Berg (2003b), and in section 4, the asymmetric Abbring and Van den Berg (2003b) model can be extended to a fully symmetric bivariate duration model in which treatment hazards may be causally affected by the past occurrence of an outcome event just like outcomes may be affected by past treatment events. This model could be used to analyze data in which both treatment and outcome times are fully observed. In this symmetric setup, any dependence in the data of the time-to-treatment hazard on past outcome events is interpreted as an effect of outcomes on future treatment decisions, and not an anticipatory effect of treatment on past outcomes. If one does not restrict the effects of outcomes on future treatment, without further

g -computation formula.

HN’s very different approach to identification allows them to incorporate anticipation effects. As in their analysis of the duration model, they assume that there is an exogenous source of independent variation of treatment decisions, independent of past outcomes. Any variation in current outcomes with variation in future treatment decisions induced by this exogenous source cannot be due to selection effects (since they explicitly control for the unobservables) and is interpreted as anticipatory effects of treatment in their framework. However, their structural analysis naturally excludes such effects (see section 5.2 below). Therefore, a natural interpretation of the ability of HN to identify anticipatory effects is that they have overidentifying restrictions that allow them to test their model and, if necessary, relax their assumptions.

In a model with uncertainty, agents act on and value *ex ante* outcomes. The model developed below in section 5.2 distinguishes *ex ante* from *ex post* outcomes. The model developed in this section cannot because, within it, it is difficult to specify the information sets on which agents act or the mechanism by which agents forecast and act on $Y(s, X, U(s))$ when they are making choices.

One justification for not making an *ex ante* – *ex post* distinction is that the agents being modeled operate under perfect foresight even though econometricians do not observe all of the information available to the agents. In this framework, the $U(s), s = 1, \dots, \bar{S}$, are an ingredient of the econometric model that accounts for the asymmetry of information between the agent and the econometrician studying the agent.

Without imposing assumptions about the functional structure of the outcome equations, it is not possible to nonparametrically identify counterfactual outcome states $Y(s, X, U(s))$ that have never been observed. Thus, in a schooling example, HN assume that analysts observe life cycle outcomes for some persons for each stopping time (level of final grade completion) and our notation reflects this.⁶² However, analysts do not observe $Y(s, X, U(s))$ for all s for anyone. A person can have only one stopping time (one completed schooling level). This observational limitation creates our evaluation problem, the “fundamental problem of causal inference”.⁶³

In addition to this problem, there is the standard selection problem that the $Y(s, X, U(s))$ are only observed for persons who stop at s and not for a random sample of the population. The selected distribution may not accurately characterize the population distribution of $Y(s, X, U(s))$ for persons selected at random. Note also that without further structure, we can only identify treatment responses within a given policy environment. In another policy environment, where the rules governing selection into treatment and/or the outcomes from treatment may be different, the same time to treatment may be associated with entirely different responses.⁶⁴ We now turn to the HN analysis of identification of outcome and treatment time distributions.

restrictions, the data on treatments occurring after the outcome event carry no information on anticipatory effects of treatment on outcomes and they face an identification problem similar to that in the asymmetric case.

⁶²In practice, analysts can only observe a portion of the life cycle after treatment. See the discussion on pooling data across samples in Cunha, Heckman, and Navarro (2005) to replace missing life cycle data.

⁶³See Holland (1986) or Gill and Robins (2001).

⁶⁴This is the problem of general equilibrium effects, and leads to violation of the policy invariance conditions. See Heckman, Lochner, and Taber (1998), Heckman, LaLonde, and Smith (1999) or Abbring and Van den Berg (2003b) for discussion of this problem.

5.1.4 Identification of Outcome and Treatment Time Distributions

We assume access to a large i.i.d. sample from the distribution of $(S, Y(S, X, U(S)), X, Z)$, where S is the stopping time, X are the variables determining outcomes and Z are the variables determining choices. We also know $\Pr(S = s \mid Z = z)$, for $s = 1, \dots, \bar{S}$, from the data. For expositional convenience, we first consider the case of scalar outcomes $Y(S, X, U(S))$. An analysis for vector $Y(S, X, U(S))$ is presented in HN and is discussed below.

Consider the analysis of continuous outcomes. HN analyze more general cases. Their results extend the analyses of Heckman and Honoré (1990), Heckman (1990) and Carneiro, Hansen, and Heckman (2003) by considering choices generated by a stopping time model. To simplify the notation in this section, assume that the scalar outcome associated with stopping at time s can be written as $Y(s) = \mu(s, X) + U(s)$, where $Y(s)$ is shorthand for $Y(s, X, U(s))$. $Y(s)$ is observed only if $D(s) = 1$ where the $D(s)$ are generated by the model analyzed in Theorem 1. Write $I(s) = \Psi(s, Z(s)) - \eta(s)$. Assume that the $\Psi(s, Z(s))$ belong to the Matzkin (1992) class of functions. We use the condensed representations I , $\Psi(Z)$, η , Y , $\mu(X)$ and U as described in the previous subsection.

Heckman and Navarro permit general stochastic dependence within the components of U , within the components of η and across the two vectors. They assume that (X, Z) are independent of (U, η) . Each component of (U, η) has a zero mean. The joint distribution of (U, η) is assumed to be absolutely continuous.

With “sufficient variation” in the components of $\Psi(Z)$, one can identify $\mu(s, X)$, $[\Psi(1, Z(1)), \dots, \Psi(s, Z(s))]$ and the joint distribution of $U(s)$ and η^s . This enables the analyst to identify average treatment effects across all stopping times, since one can extract $E(Y(s) - Y(s') \mid X = x)$ from the marginal distributions of $Y(s)$, $s = 1, \dots, \bar{S}$.

Theorem 2. Write $\Psi^s(Z) = (\Psi(1, Z(1)), \dots, \Psi(s, Z(s)))$. Assume in addition to the conditions in Theorem 1 that

- (i) $E[U(s)] = 0$. $(U(s), \eta^s)$ are continuous random variables with support $\text{Supp}(U(s)) \times \text{Supp}(\eta^s)$ with upper and lower limits $(\bar{U}(s), \bar{\eta}^s)$ and $(\underline{U}(s), \underline{\eta}^s)$, respectively, $s = 1, \dots, \bar{S}$. These conditions hold for each component of each subvector. The joint system is thus variation free for each component with respect to every other component.
- (ii) $(U(s), \eta^s) \perp\!\!\!\perp (X, Z)$, $s = 1, \dots, \bar{S}$ (independence).
- (iii) $\mu(s, X)$ is a continuous function, $s = 1, \dots, \bar{S}$.
- (iv) $\text{Supp}(\Psi(Z), X) = \text{Supp}(\Psi(Z)) \times \text{Supp}(X)$.

Then one can identify $\mu(s, X)$, $\Psi^s(Z)$, $F_{\eta^s, U(s)}$, $s = 1, \dots, \bar{S}$, where $\Psi(Z)$ is identified over the support admitted by condition (ii) of Theorem 1.

Proof. See Abbring and Heckman (2007), Appendix C. □

The proof in Abbring and Heckman (2007, Appendix C) covers the case of vector $Y(s, X, U(s))$ where each component is a continuous random variable. Appendix D of Abbring and Heckman (2007) states and proves a more general theorem for age-specific outcomes $Y(t, s, X, U(t, s))$, $t = 1, \dots, \bar{T}$, where Y can be a vector of continuous and discrete outcomes. In particular, HN can identify age-specific earnings flows associated with multiple sources of income.

Theorem 2 does not identify the joint distribution of $Y(1), \dots, Y(\bar{S})$ because analysts observe only one of these outcomes for any person. Observe that exclusion restrictions in the arguments of the choice of treatment equation are not required to identify the counterfactuals. What is required is independent variation of arguments which might be achieved by exclusion conditions but can be obtained by other functional restrictions (see HN, Corollary 1, for example). One can identify the $\mu(s, X)$ (up to constants) without the limit set argument. Thus, one can identify certain features of the model without using the limit set argument. See HN.

As a by-product of Theorem 2, one can construct various counterfactual distributions of $Y(s)$ for agents with index crossing histories such that $D(s) = 0$ (that is, for whom $Y(s)$ is not observed). Define $B(s) = \mathbf{1}[I(s) \geq 0]$, $B^s = (B(1), \dots, B(s))$, and let b^s denote a vector of possible values of B^s . $D(s)$ was defined as $B(s)$ if $B^{s-1} = (0)$ and 0 otherwise. Theorem 2 gives conditions under which the counterfactual distribution of $Y(s)$ for those with $D(s') = 1$, $s' \neq s$, can be constructed. More generally, it can be used to construct

$$\Pr(Y(s) \leq y(s) \mid B^{s'} = b^{s'}, X = x, Z = z)$$

for all of the $2^{s'}$ possible sequences $b^{s'}$ of $B^{s'}$ outcomes up to $s' \leq s$. If $b^{s'}$ equals a sequence of $s' - 1$ zeros followed by a one, then $B^{s'} = b^{s'}$ corresponds to $D(s') = 1$. The event $B^{s'} = (0)$ corresponds to $D(s') = (0)$, i.e., $S > s'$. For all other sequences $b^{s'}$, $B^{s'} = b^{s'}$ defines a subpopulation of the agents with $D(s'') = 1$ for some $s'' < s'$ and multiple index crossings. For example, $B^{s'} = (0, 1, 0)$ corresponds to $D(2) = 1$ and $I(3) < 0$. This defines a subpopulation that takes treatment at time 2, but that would not take treatment at time 3 if it would not have taken treatment at time 2.⁶⁵ It is tempting to interpret such sequences with multiple crossings as corresponding to multiple entry into and exit from treatment. However, this is inconsistent with the stopping time model (11), and would require extension of the model to deal with recurrent treatment. Whether a threshold-crossing model corresponds to a structural model of treatment choice is yet another issue, which is taken up in the next section and is also addressed in Cunha, Heckman, and Navarro (2007).

The counterfactuals that are identified by fixing $D(s') = 1$ for different treatment times s' in the general model of HN have an asymmetric aspect. HN can generate $Y(s)$ distributions for persons who are treated at s or before. Without further structure, they cannot generate the distributions of these random variables for people who receive treatment at times after s .

The source of this asymmetry is the generality of duration model (11). At each stopping time s , HN acquire a new random variable $\eta(s)$ which can have arbitrary dependence with $Y(s)$ and $Y(s')$ for all s and s' . From Theorem 2, HN can identify the dependence between $\eta(s')$ and $Y(s)$ if $s' \leq s$. They cannot identify the dependence between $\eta(s')$ and $Y(s)$ for $s' > s$ without imposing further

⁶⁵Cunha, Heckman, and Navarro (2007) develop an ordered choice model with stochastic thresholds.

structure on the unobservables.⁶⁶ Thus, one can identify the distribution of college outcomes for high school graduates who do not go on to college and can compare these to outcomes for high school graduates, so they can identify the parameter “treatment on the untreated.” However, one cannot identify the distribution of high school outcomes for college graduates (and hence treatment on the treated parameters) without imposing further structure.⁶⁷ Since one can identify the marginal distributions under the conditions of Theorem 2, one can identify pairwise average treatment effects for all s, s' .

It is interesting to contrast the model identified by Theorem 2 with a conventional static multinomial discrete-choice model with an associated system of counterfactuals, as presented in Heckman and Vytlacil (2007a, Appendix B) and analyzed in Abbring and Heckman (2007, section 2). Using standard tools, it is possible to establish semiparametric identification of the conventional static model of discrete choice joined with counterfactuals and to identify all of the standard mean counterfactuals. For that model there is a fixed set of unobservables governing all choices of states. Thus the analyst does not acquire new unobservables associated with each stopping time as occurs in a dynamic model. In a dynamic model, selection effects for $Y(s)$ depend on the unobservables up to s but not later innovations. Selection effects in a static discrete-choice model depend on a fixed set of unobservables for all outcomes. With suitable normalizations, HN identify the joint distributions of choices and associated outcomes without the difficulties, just noted, that appear in the reduced-form dynamic model. HN develop models for discrete outcomes including duration models.

5.1.5 Using Factor Models to Identify Joint Distributions of Counterfactuals

From Theorem 2 and its generalizations reported in HN, one can identify joint distributions of outcomes for each treatment time s and the index generating treatment times. One cannot identify the joint distributions of outcomes across treatment times. Moreover, as just discussed, one cannot, in general, identify treatment on the treated parameters.

Aakvik, Heckman, and Vytlacil (2005) and Carneiro, Hansen, and Heckman (2003) show how to use factor models to identify the joint distributions across treatment times and recover the standard treatment parameters. HN use their approach to identify the joint distribution of $Y = (Y(1), \dots, Y(\bar{S}))$.

The basic idea underlying this approach is to use joint distributions for outcomes measured at each treatment time s along with the choice index to construct the joint distribution of outcomes across treatment choices. To illustrate how to implement this intuition, suppose that we augment Theorem 2 by appealing to Theorem 2 in Carneiro, Hansen, and Heckman (2003) to identify the joint distribution of the vector of outcomes at each stopping time along with $I^s = (I(1), \dots, I(s))$

⁶⁶One possible structure is a factor model which is applied to this problem in the next section.

⁶⁷In the schooling example, one can identify treatment on the treated for the final category \bar{S} since $D^{\bar{S}-1} = (0)$ implies $D(\bar{S}) = 1$. Thus at stage $\bar{S} - 1$, one can identify the distribution of $Y(\bar{S} - 1)$ for persons for whom $D(0) = 0, \dots, D(\bar{S} - 1) = 0, D(\bar{S}) = 1$. Hence, if college is the terminal state, and high school the state preceding college, one can identify the distribution of high school outcomes for college graduates.

for each s . For each s , we may write

$$\begin{aligned} Y(t, s, X, U(t, s)) &= \mu(t, s, X) + U(t, s), \quad t = 1, \dots, \bar{T} \\ I(s) &= \Psi(s, Z(s)) - \eta(s). \end{aligned}$$

The scale of $\Psi(s, Z(s))$ is determined from the Matzkin (1992) conditions. If we specify the Matzkin functions only up to scale, we determine the functions up to scale and make a normalization. From Theorem 2, we can identify the joint distribution of $(\eta(1), \dots, \eta(s), U(1, s), \dots, U(\bar{T}, s))$.

To review these concepts and their application to the model discussed in this section, suppose that we adopt a one-factor model where θ is the factor. It has mean zero. The errors can be represented by

$$\begin{aligned} \eta(s) &= \varphi_s \theta + \varepsilon_{\eta(s)} \\ U(t, s) &= \alpha_{t,s} \theta + \varepsilon_{t,s}, \quad t = 1, \dots, \bar{T}, \quad s = 1, \dots, \bar{S}. \end{aligned}$$

The θ are independent of all of the $\varepsilon_{\eta(s)}$, $\varepsilon_{t,s}$ and the ε 's are mutually independent mean zero disturbances. The φ_s and $\alpha_{t,s}$ are factor loadings. Since θ is an unobservable, its scale is unknown. One can set the scale of θ by normalizing one factor loading, say $\alpha_{\bar{T}, \bar{S}} = 1$. From the joint distribution of $(\eta^s, U(s))$, one can identify σ_θ^2 , $\alpha_{t,s}, \varphi_s, t = 1, \dots, \bar{T}$, for $s = 1, \dots, \bar{S}$, using the arguments presented in, e.g., Abbring and Heckman (2007, section 2.8). A sufficient condition is $\bar{T} \geq 3$, but this ignores possible additional information from cross-system restrictions. From this information, one can form for $t \neq t'$ or $s \neq s''$ or both,

$$\text{Cov}(U(t, s), U(t', s'')) = \alpha_{t,s} \alpha_{t',s''} \sigma_\theta^2,$$

even though the analyst does not observe outcomes for the same person at two different stopping times. In fact, one can construct the joint distribution of $(U, \eta) = (U(1), \dots, U(\bar{S}), \eta)$. From this joint distribution, one can recover the standard mean treatment effects as well as the joint distributions of the potential outcomes. One can determine the percentage of participants at treatment time s who benefit from participation compared to what their outcomes would be at other treatment times. One can perform a parallel analysis for models for discrete outcomes and durations. The analysis can be generalized to multiple factors. Conventional factor analysis assumes that the unobservables are normally distributed. Carneiro, Hansen, and Heckman (2003) establish nonparametric identifiability of the θ 's and the ε 's and their analysis of nonparametric identifiability applies here.

Theorem 2, strictly applied, actually produces only one scalar outcome along with one or more choices for each stopping time.⁶⁸ If vector outcomes are not available, access to a measurement system M that assumes the same values for each stopping time can substitute for the need for vector outcomes for Y . Let M_j be the j^{th} component of this measurement system. Write

$$M_j = \mu_{j,M}(X) + U_{j,M}, \quad j = 1, \dots, J,$$

where $U_{j,M}$ are mean zero and independent of X .

⁶⁸HN and Abbring and Heckman (2007) analyze the vector-outcome case.

Suppose that the $U_{j,M}$ have a one-factor structure so $U_{j,M} = \alpha_{j,M}\theta + \varepsilon_{j,M}$, $j = 1, \dots, J$, where the $\varepsilon_{j,M}$ are mean zero, mutually independent random variables, independent of the θ . Adjoining these measurements to the one outcome measure $Y(s)$ can substitute for the measurements of $Y(t, s)$ used in the previous example. In an analysis of schooling, the M_j can be test scores that depend on ability θ . Ability is assumed to affect outcomes $Y(s)$ and the choice of treatment times indices.

The factor models implement a matching on unobservables assumption, $\{Y(s)\}_{s=1}^{\bar{S}} \perp\!\!\!\perp S \mid X, Z, \theta$. HN allow for the θ to be unobserved variables and present conditions under which their distributions can be identified.

5.1.6 Summary of the Reduced-Form Model

A limitation of the reduced-form approach pursued in this section is that, because the underlying model of choice is not clearly specified, it is not possible without further structure to form, or even define, the marginal treatment effect analyzed in Heckman and Vytlačil (1999, 2001, 2005, 2007a,b) or Heckman, Urzua, and Vytlačil (2006). The absence of well defined choice equations is problematic for the models analyzed thus far in this section of our chapter, although it is typical of many statistical treatment effect analyses.⁶⁹ In this framework, it is not possible to distinguish objective outcomes from subjective evaluations of outcomes, and to distinguish *ex ante* from *ex post* outcomes. Another limitation of this analysis is its strong reliance on large support conditions on the regressors coupled with independence assumptions. Independence can be relaxed following Lewbel (2000) and Honoré and Lewbel (2002). The large support assumption plays a fundamental role here and throughout the entire evaluation literature.

HN develop an explicit economic model for dynamic treatment effects that allows analysts to make these and other distinctions. They extend the analysis presented in this subsection to a more precisely formulated economic model. They explicitly allow for agent updating of information sets. A well posed economic model enables economists to evaluate policies in one environment and accurately project them to new environments as well as to accurately forecast new policies never previously experienced. We now turn to an analysis of a more fully articulated structural econometric model.

5.2 A Sequential Structural Model with Option Values

This section analyzes the identifiability of a structural sequential optimal stopping time model. HN use ingredients assembled in the previous sections to build an economically interpretable framework for analyzing dynamic treatment effects. For specificity, HN focus on a schooling model with associated earnings outcomes that is motivated by the research of Keane and Wolpin (1997) and Eckstein and Wolpin (1999). They explicitly model costs and build a dynamic version of a Roy model. We briefly survey the literature on dynamic discrete choice in section 5.5 below.

In the model of this section, it is possible to interpret the literature on dynamic treatment effects

⁶⁹Heckman (2005) and Heckman and Vytlačil (2007a,b) point out that one distinctive feature of the economic approach to program evaluation is the use of choice theory to define parameters and evaluate alternative estimators.

within the context of an economic model; to allow for earnings while in school as well as grade-specific tuition costs; to separately identify returns and costs; to distinguish private evaluations from “objective” *ex ante* and *ex post* outcomes and to identify persons at various margins of choice. In the context of medical economics, HN consider how to identify the pain and suffering associated with a treatment as well as the distribution of benefits from the intervention. They also model how anticipations about potential future outcomes associated with various choices evolve over the life cycle as sequential treatment choices are made.

In contrast to the analysis of section 5.1, the identification proof for their dynamic choice model works in reverse starting from the last period and sequentially proceeding backward. This approach is required by the forward-looking nature of dynamic choice analysis and makes an interesting contrast with the analysis of identification for the reduced-form models which proceeds forward from initial period values.

HN use limit set arguments to identify the parameters of outcome and measurement systems for each stopping time $s = 1, \dots, \bar{S}$, including means and joint distributions of unobservables. These systems are identified without invoking any special assumptions about the structure of model unobservables. When they invoke factor structure assumptions for the unobservables, they identify the factor loadings associated with the measurements (as defined in section 5.1.5) and outcomes. They also nonparametrically identify the distributions of the factors and the distributions of the innovations to the factors. With the joint distributions of outcomes and measurements in hand for each treatment time, HN can identify cost (and preference) information from choice equations that depend on outcomes and costs (preferences). HN can also identify joint distributions of outcomes across stopping times. Thus, they can identify the proportion of people who benefit from treatment. Their analysis generalizes the one shot decision models of Cunha and Heckman (2007, 2008); Cunha, Heckman, and Navarro (2005, 2006) to a sequential setting.

All agents start with one year of schooling at age 1 and then sequentially choose, at each subsequent age, whether to continue for another year in school. New information arrives at each age. One of the benefits of staying in school is the arrival of new information about returns. Each year of schooling takes one year of age to complete. There is no grade repetition. Once persons leave school, they never return.⁷⁰ As a consequence, an agent’s schooling level equals her age up to the time $S \leq \bar{S}$ she leaves school. After that, ageing continues up to age $\bar{T} \geq \bar{S}$, but schooling does not. We again denote $D(s) = \mathbf{1}(S = s)$ for all $s \in \{1, \dots, \bar{S}\}$. Let $\delta(t) = 1$ if a person has left school at or before age t ; $\delta(t) = 0$ if a person is still in school.

A person’s earnings at age t depend on her *current* schooling level s , and whether she has left school on or before age t ($\delta(t) = 1$) or not ($\delta(t) = 0$). Thus,

$$Y(t, s, \delta(t), X) = \mu(t, s, \delta(t), X) + U(t, s, \delta(t)). \quad (12)$$

Note that $Y(t, s, 0, X)$ is only meaningfully defined if $s = t$, in which case it denotes the earnings of a person as a student at age and schooling level s . More precisely, $Y(s, s, 0, X)$ denotes the earnings

⁷⁰It would be better to derive such stopping behavior as a feature of a more general model with possible recurrence of states. Cunha, Heckman, and Navarro (2007) develop general conditions under which it is optimal to stop and never return.

of an individual with characteristics X who is still enrolled in school at age and schooling level s and goes on to complete at least $s+1$ years of schooling. The fact that earnings in school depend only on the current schooling level, and not on the final schooling level obtained, reflects the no-anticipation condition (NA). $U(t, s, \delta(t))$ is a mean zero shock that is unobserved by the econometrician but may, or may not, be observed by the agent. $Y(t, s, 1, X)$ is meaningfully defined only if $s \leq t$, in which case it denotes the earnings at age t of an agent who has decided to stop schooling at s .

The direct cost of remaining enrolled in school at age and schooling level s is

$$C(s, X, Z(s)) = \Phi(s, X, Z(s)) + W(s)$$

where X and $Z(s)$ are vectors of observed characteristics (from the point of view of the econometrician) that affect costs at schooling level s , and $W(s)$ are mean zero shocks that are unobserved by the econometrician that may or may not be observed by the agent. Costs are paid in the period before schooling is undertaken. The agent is assumed to know the costs of making schooling decisions at each transition. The agent is also assumed to know the X and $Z = (Z(1), \dots, Z(\bar{S}-1))$ from age 1.⁷¹

The optimal schooling decision involves comparisons of the value of continuing in school for another year and the value of leaving school forever at each age and schooling level $s \in \{1, \dots, \bar{S}-1\}$. We can solve for these values, and the optimal schooling decision, by backward recursion.

The agent's expected reward of stopping schooling forever at level and age s (i.e., receiving treatment s) is given by the expected present value of her remaining lifetime earnings:

$$R(s, I_s) = E \left(\sum_{j=0}^{\bar{T}-s} \left(\frac{1}{1+r} \right)^j Y(s+j, s, 1, X) \mid I_s \right), \quad (13)$$

where I_s are the state variables generating the age- s -specific information set \mathcal{I}_s .⁷² They include the schooling level attained at age s , the covariates X and Z , as well as all other variables known to the agent and used in forecasting future variables. Assume a fixed, nonstochastic, interest rate r .⁷³ The continuation value at age and schooling level s given information I_s is denoted by $K(s, I_s)$.

At $\bar{S}-1$, when an individual decides whether to stop or continue on to \bar{S} , the expected reward from remaining enrolled and continuing to \bar{S} (i.e., the continuation value) is the earnings while in school less costs plus the expected discounted future return that arises from completing \bar{S} years of schooling:

$$\begin{aligned} K(\bar{S}-1, I_{\bar{S}-1}) &= Y(\bar{S}-1, \bar{S}-1, 0, X) - C(\bar{S}-1, X, Z(\bar{S}-1)) \\ &\quad + \frac{1}{1+r} E(R(\bar{S}, I_{\bar{S}}) \mid I_{\bar{S}-1}) \end{aligned}$$

⁷¹These assumptions can be relaxed and are made for convenience. See Carneiro, Hansen, and Heckman (2003), Cunha, Heckman, and Navarro (2005) and Cunha and Heckman (2007) for a discussion of selecting variables in the agent's information set.

⁷²We only consider the agent's information set here, and drop the subscript A for notational convenience.

⁷³This assumption is relaxed in HN who present conditions under which r can be identified.

where $C(\bar{S} - 1, X, Z(\bar{S} - 1))$ is the direct cost of schooling for the transition to \bar{S} . This expression embodies the assumption that each year of school takes one year of age. $I_{\bar{S}-1}$ incorporates all of the information known to the agent.

The value of being in school just before deciding on continuation at age and schooling level $\bar{S} - 1$ is the larger of the two expected rewards that arise from stopping at $\bar{S} - 1$ or continuing one more period to \bar{S} :

$$V(\bar{S} - 1, I_{\bar{S}-1}) = \max \{R(\bar{S} - 1, I_{\bar{S}-1}), K(\bar{S} - 1, I_{\bar{S}-1})\}.$$

More generally, at age and schooling level s this value is

$$\begin{aligned} V(s, I_s) &= \max \{R(s, I_s), K(s, I_s)\} \\ &= \max \left\{ R(s, I_s), \left(\frac{Y(s, s, 0, X) - C(s, X, Z(s))}{1 + r} + E(V(s + 1, I_{s+1}) | I_s) \right) \right\}.^{74} \end{aligned}$$

Following the exposition of the reduced-form decision rule in section 5.1, define the decision rule in terms of a first passage of the “index” $R(s, I_s) - K(s, I_s)$,

$$D(s) = \mathbf{1}[R(s, I_s) - K(s, I_s) \geq 0, R(s - 1, I_{s-1}) - K(s - 1, I_{s-1}) < 0, \dots, R(1, I_1) - K(1, I_1) < 0].$$

An individual stops at the schooling level at the first age where this index becomes positive. From data on stopping times, one can nonparametrically identify the conditional probability of stopping at s ,

$$\Pr(S = s | X, Z) = \Pr \left(\begin{array}{l} R(s, I_s) - K(s, I_s) \geq 0, \\ R(s - 1, I_{s-1}) - K(s - 1, I_{s-1}) < 0, \dots, \\ R(1, I_1) - K(1, I_1) < 0 \end{array} \middle| X, Z \right).$$

HN use factor structure models based on the θ introduced in section 5.1 to define the information updating structure. Agents learn about different components of θ as they evolve through life. The HN assumptions allow for the possibility that agents may know some or all the elements of θ at a given age t regardless of whether or not they determine earnings at or before age t . Once known,

⁷⁴This model allows no recall and is clearly a simplification of a more general model of schooling with option values. Instead of imposing the requirement that once a student drops out the student never returns, it would be useful to derive this property as a feature of the economic environment and the characteristics of individuals. Cunha, Heckman, and Navarro (2007) develop such conditions. In a more general model, different persons could drop out and return to school at different times as information sets are revised. This would create further option value beyond the option value developed in the text that arises from the possibility that persons who attain a given schooling level can attend the next schooling level in any future period. Implicit in this analysis of option values is the additional assumption that persons must work at the highest level of education for which they are trained. An alternative model allows individuals to work each period at the highest wage across all levels of schooling that they have attained. Such a model may be too extreme because it ignores the costs of switching jobs, especially at the higher educational levels where there may be a lot of job-specific human capital for each schooling level. A model with these additional features is presented in Heckman, Urzua, and Yates (2007).

they are not forgotten. As agents accumulate information, they revise their forecasts of their future earnings prospects at subsequent stages of the decision process. This affects their decision rules and subsequent choices. Thus HN allow for learning which can affect both pretreatment outcomes and posttreatment outcomes.^{75,76} All dynamic discrete-choice models make some assumptions about the updating of information and any rigorous identification analysis of this class of models must test among competing specifications of information updating.

Variables unknown to the agent are integrated out by the agent in forming expectations over future outcomes. Variables known to the agent are treated as constants by the agents. They are integrated out by the econometrician to control for heterogeneity. These are separate operations except for special cases. In general, the econometrician knows less than what the agent knows. The econometrician seeks to identify the distributions of the variables in the agent information sets that are used by the agents to form their expectations as well as the distributions of variables known to the agent and treated as certain quantities by the agent but not known by the econometrician. Determining which elements belong in the agent's information set can be done using the methods expositied in Cunha, Heckman, and Navarro (2005) and Cunha and Heckman (2007) who consider testing what components of X, Z, ε as well as θ are in the agent's information set. We briefly discuss this issue at the end of the next section.⁷⁷ HN establish semiparametric identification of the model assuming a given information structure. Determining the appropriate information structure facing the agent and its evolution is an essential aspect of identifying any dynamic discrete-choice model.

Observe that agents with the same information variables I_t at age t have the same expectations of future returns, and the same continuation and stopping values. They make the same investment choices. Persons with the same *ex ante* reward, state and preference variables have the same *ex ante* distributions of stopping times. *Ex post*, stopping times may differ among agents with identical *ex ante* information. Controlling for I_t , future realizations of stopping times do not affect past rewards. This rules out the problem that the future can cause the past, which may happen in HN's reduced-form model. It enforces the (NA) condition of Abbring and Van den Berg. Failure to accurately model I_t produces failure of (NA).

HN establish semiparametric identification of their model without period-by-period exclusion restrictions. Their analysis extends Theorems 1 and 2 to an explicit choice-theoretic setting. They use limit set arguments to identify the joint distributions of earnings (for each treatment time

⁷⁵This type of learning about unobservables can be captured by HN's reduced-form model, but not by Abbring and Van den Berg's (2003b) single-spell mixed proportional hazards model. Their model does not allow for time-varying unobservables. Abbring and Van den Berg develop a multiple-spell model that allows for time-varying unobservables. Moreover, their nonparametric discussion of (NA) and randomization does not exclude the sequential revelation to the agent of a finite number of unobserved factors although they do not systematically develop such a model.

⁷⁶It is fruitful to distinguish models with exogenous arrival of information (so that information arrives at each age t independent of any actions taken by the agent) from information that arrives as a result of choices by the agent. The HN model is in the first class. The models of Miller (1984) or Pakes (1986) are in the second class. See our discussion in section 5.5.

⁷⁷The HN model of learning is clearly very barebones. Information arrives exogenously across ages. In the factor model, all agents who advance to a stage get information about additional factors at that stage of their life cycles but the realizations of the factors may differ across persons.

s across t) and any associated measurements that do not depend on the stopping time chosen. For each stopping time, they construct the means of earnings outcomes at each age and of the measurements and the joint distributions of the unobservables for earnings and measurements. Factor analyzing the joint distributions of the unobservables, under conditions specified in Carneiro, Hansen, and Heckman (2003), they identify the factor loadings, and nonparametrically identify the distributions of the factors and the independent components of the error terms in the earnings and measurement equations. Armed with this knowledge, they use choice data to identify the distribution of the components of the cost functions that are not directly observed. They construct the joint distributions of outcomes across stopping times. They also present conditions under which the interest rate r is identified.

In their model, analysts can distinguish period by period *ex ante* expected returns from *ex post* realizations by applying the analysis of Cunha, Heckman, and Navarro (2005). See the survey in Heckman, Lochner, and Todd (2006) and section 2 of Abbring and Heckman (2007) for discussions of this approach. Because they link choices to outcomes through the factor structure assumption, they can also distinguish *ex ante* preference or cost parameters from their *ex post* realizations. *Ex ante*, agents may not know some components of θ . *Ex post*, they do. All of the information about future rewards and returns is embodied in the information set \mathcal{I}_t . Unless the time of treatment is known with perfect certainty, it cannot cause outcomes prior to its realization.

The analysis of HN is predicated on specification of agent information sets. These information sets should be carefully distinguished from those of the econometrician. Cunha, Heckman, and Navarro (2005) present methods for determining which components of future outcomes are in the information sets of agents at each age, \mathcal{I}_t . If there are components unknown to the agent at age t , under rational expectations, agents form their value functions used to make schooling choices by integrating out the unknown components using the distributions of the variables in their information sets. Components that are known to the agent are treated as constants by the individual in forming the value function but as unknown variables by the econometrician and their distribution is estimated. The true information set of the agent is determined from the set of possible specifications of the information sets of agents by picking the specification that best fits the data on choices and outcomes penalizing for parameter estimation. If neither the agent nor the econometrician knows a variable, the econometrician identifies the determinants of the distribution of the unknown variables that is used by the agent to form expectations. If the agent knows some variables, but the econometrician does not, the econometrician seeks to identify the distribution of the variables, but the agent treats the variables as known constants.

HN can identify all of the treatment parameters including the pairwise average treatment effect (ATE), the marginal treatment effect (MTE) for each transition (obtained by finding mean outcomes for individuals indifferent between transitions), all of the treatment on the treated and treatment on the untreated parameters and the population distribution of treatment effects by applying the analysis of Carneiro, Hansen, and Heckman (2003) and Cunha, Heckman, and Navarro (2005) to this model. Their analysis can be generalized to cover the case where there are vectors of

contemporaneous outcome measures for different stopping times. See HN for proofs and details.⁷⁸

5.3 Identification at Infinity

Heckman and Navarro (2007), and many other researchers, rely on identification at infinity to obtain their main identification results. Identification at infinity is required to identify the average treatment effect (ATE) using IV and control function methods and in the reduced-form discrete-time models developed in the previous subsections. While this approach is controversial, it is also testable. In any sample, one can plot the distributions of the probability of each state (exit time) to determine if the identification conditions are satisfied in any sample. Figure 1, presented by HN from the research of Heckman, Stixrud, and Urzua (2006), shows such plots for a six-state static schooling model that they estimate. To identify the marginal outcome distributions for each state, the support of the state probabilities should be the full unit interval. The identification at infinity condition is clearly not satisfied in their data.⁷⁹ Only the empirical distribution of the state probability of graduating from a four-year college comes even close to covering the full unit interval. Thus, their empirical results rely on parametric assumptions, and ATE and the marginal distributions of outcomes are nonparametrically nonidentified in their data without invoking additional structure.

5.4 Comparing Reduced-Form and Structural Models

The reduced-form model analyzed in section 5.1 is typical of many reduced-form statistical approaches within which it is difficult to make important conceptual distinctions. Because agent choice equations are not modeled explicitly, it is hard to use such frameworks to formally analyze the decision makers' expectations, costs of treatment, the arrival of information, the content of agent information sets and the consequences of the arrival of information for decisions regarding time to treatment as well as outcomes. Key behavioral assumptions are buried in statistical assumptions. It is difficult to distinguish *ex post* from *ex ante* valuations of outcomes in the reduced-form models. Cunha, Heckman, and Navarro (2005), Carneiro, Hansen, and Heckman (2003) and Cunha and Heckman (2007, 2008) present analyses that distinguish *ex ante* anticipations from *ex post* realizations.⁸⁰ In reduced-form models, it is difficult to make the distinction between private evaluations and preferences (e.g., "costs" as defined in this section) from objective outcomes (the Y variables).

Statistical and reduced-form econometric approaches to analyzing dynamic counterfactuals appeal to uncertainty to motivate the stochastic structure of models. They do not explicitly characterize how agents respond to uncertainty or make treatment choices based on the arrival of new information (see Robins, 1989, 1997, Lok, 2007, Gill and Robins, 2001, Abbring and Van den Berg, 2003b, and Van der Laan and Robins, 2003). The structural approach surveyed in section 5.2 and developed by HN allows for a clear treatment of the arrival of information, agent expectations,

⁷⁸The same limitations regarding independence assumptions between the regressors and errors discussed in the analysis of reduced forms apply to the structural model.

⁷⁹One can always argue that they are satisfied in an infinite sample that has not yet been realized. That statement has no empirical content.

⁸⁰See the summary of this literature in Heckman, Lochner, and Todd (2006).

and the effects of new information on choice and its consequences. In an environment of imperfect certainty about the future, it rules out the possibility of the future causing the past once the effects of agent information are controlled for.

The structural model developed by HN allows agents to learn about new factors (components of θ) as they proceed sequentially through their life cycles. It also allows agents to learn about other components of the model (see Cunha, Heckman, and Navarro, 2005). Agent anticipations of when they will stop and the consequences of alternative stopping times can be sequentially revised. Agent anticipated payoffs and stopping times are sequentially revised as new information becomes available. The mechanism by which agents revise their anticipations is modeled and identified. See Cunha, Heckman, and Navarro (2005, 2006), Cunha and Heckman (2007, 2008) and Abbring and Heckman (2007) for further discussion of these issues and Heckman, Lochner, and Todd (2006) for a partial survey of recent developments in the literature.

The clearest interpretation of the models in the statistical literature on dynamic treatment effects is as *ex post* selection-corrected analyses of distributions of events that have occurred. In a model of perfect certainty, where *ex post* and *ex ante* choices and outcomes are identical, the reduced-form approach can be interpreted as approximating clearly specified choice models. In a more general analysis with information arrival and agent updating of information sets, the nature of the approximation is less clear cut. Thus, the current reduced-form literature is unclear as to which agent decision-making processes and information arrival assumptions justify the conditional sequential randomization assumptions widely used in the dynamic treatment effect literature (see, e.g., Gill and Robins, 2001; Lechner and Miquel, 2002; Lok, 2007; Robins, 1989, 1997; Van der Laan and Robins, 2003). Section 3.2.2 provides some insight by highlighting the connection to the conditional-independence assumption often employed in the structural dynamic discrete-choice literature (see Rust, 1987; and the survey in Rust, 1994). Reduced-form approaches are not clear about the source of the unobservables and their relationship with conditioning variables. It would be a valuable exercise to exhibit which structural models are approximated by various reduced-form models. In the structural analysis, this specification emerges as part of the analysis, as our discussion of the stochastic properties of the unobservables presented in the preceding section makes clear.

The HN analysis of both structural and reduced-form models relies heavily on limit set arguments. They solve the selection problem in limit sets. The dynamic matching models of Gill and Robins (2001) and Lok (2007) solve the selection problem by invoking recursive conditional independence assumptions. In the context of the models of HN, they assume that the econometrician knows the θ or can eliminate the effect of θ on estimates of the model by conditioning on a suitable set of variables. The HN analysis entertains the possibility that analysts know less than the agents they study. It allows for some of the variables that would make matching valid to be unobservable. As we have noted in early subsections, versions of recursive conditional independence assumptions are also used in the dynamic discrete-choice literature (see the survey in Rust, 1994). The HN factor models allow analysts to construct the joint distribution of outcomes across stopping times. This feature is missing from the statistical treatment effect literature.

Both HN's structural and reduced-form models of treatment choice are stopping time models.

Neither model allows for multiple entry into and exit from treatment, even though agents in these models would like to reverse their treatment decisions for some realizations of their index if this was not too costly (or, in the case of the reduced-form model, if the index thresholds for returning would not be too low).⁸¹ Cunha, Heckman, and Navarro (2007) derive conditions on structural stopping models from a more basic model that entertains the possibility of return from dropout states but which nonetheless exhibits the stopping time property. The HN identification strategy relies on the nonrecurrent nature of treatment. Their identification strategy of using limit sets can be applied to a recurrent model provided that analysts confine attention to subsets of (X, Z) such that in those subsets the probability of recurrence is zero.

5.5 A Short Survey of Dynamic Discrete-Choice Models

Rust (1994) presents a widely cited nonparametric nonidentification theorem for dynamic discrete-choice models. It is important to note the restrictive nature of his negative results. He analyzes a recurrent state infinite horizon model in a stationary environment. He does not use any exclusion restrictions or cross outcome-choice restrictions. He uses a general utility function. He places no restrictions on period-specific utility functions such as concavity or linearity nor does he specify restrictions connecting preferences and outcomes. One can break Rust's nonidentification result with additional information.

Magnac and Thesmar (2002) present an extended comment on Rust's analysis including positive results for identification when the econometrician knows the distributions of unobservables, assumes that unobservables enter period-specific utility functions in an additively separable way and is willing to specify functional forms of utility functions or other ingredients of the model, as do Pakes (1986), Keane and Wolpin (1997), Eckstein and Wolpin (1999), and Hotz and Miller (1988, 1993). Magnac and Thesmar (2002) also consider the case where one state (choice) is absorbing (as do Hotz and Miller, 1993) and where the value functions are known at the terminal age (\bar{T}) (as do Keane and Wolpin, 1997 and Belzil and Hansen, 2002). In HN, each treatment time is an absorbing state. In a separate analysis, Magnac and Thesmar consider the case where unobservables from the point of view of the econometrician are correlated over time (or age t) and choices (s) under the assumption that the distribution of the unobservables is known. They also consider the case where exclusion restrictions are available. Throughout their analysis, they maintain that the distribution of the unobservables is known both by the agent and the econometrician.

HN provide semiparametric identification of a finite-horizon finite-state model with an absorbing state with semiparametric specifications of reward and cost functions.⁸² Given that rewards are in value units, the scale of their utility function is fixed as they also are in models of profit-maximizing firms. Choices are not invariant to arbitrary affine transformations so that one source

⁸¹Recall that treatment occurs if the index turns positive. If there are costs to reversing this decision, agents would only reverse their decision if the index falls below some negative threshold. The stopping time assumption is equivalent to the assumption that the costs of reversal are prohibitively large, or that the corresponding threshold is at the lower end of the support of the index.

⁸²Although their main theorems are for additively separable reward and cost functions, it appears that additive separability can be relaxed using the analysis of Matzkin (2003).

of nonidentifiability in Rust's analysis is eliminated. They can identify the error distributions non-parametrically given their factor structure. They do not have to assume either the functional form of the unobservables or knowledge of the entire distribution of unobservables.

HN present a fully specified structural model of choices and outcomes motivated by, but not identical to, the analyses of Keane and Wolpin (1994, 1997) and Eckstein and Wolpin (1999). In their setups, outcome and cost functions are parametrically specified. Their states are recurrent while those of HN are absorbing. In the HN model, once an agent drops out of school, the agent does not return. In the Keane-Wolpin model, an agent who drops out can return. Keane and Wolpin do not establish identification of their model, whereas HN establish semiparametric identification of their model. They analyze models with more general times series processes for unobservables. In both the HN and Keane-Wolpin frameworks, agents learn about unobservables. In the Keane-Wolpin framework, such learning is about temporally independent shocks that do not affect agent expectations about returns relevant to possible future choices. The information just affects the opportunity costs of current choices. In the HN framework, learning affects agent expectations about future returns as well as current opportunity costs.

The HN model extends previous work by Carneiro, Hansen, and Heckman (2003) and Cunha and Heckman (2007, 2008); Cunha, Heckman, and Navarro (2005, 2006) by considering explicit multiperiod dynamic models with information updating. They consider one-shot decision models with information updating and associated outcomes.

Their analysis is related to that of Taber (2000). Like Cameron and Heckman (1998), both HN and Taber use identification-in-the-limit arguments.⁸³ Taber considers identification of a two period model with a general utility function, whereas in section 5.2, we discuss how HN consider identification of a specific form of the utility function (an earnings function) for a multiperiod maximization problem. As in HN, Taber allows for the sequential arrival of information. His analysis is based on conventional exclusion restrictions, but the analysis of HN is not. They use outcome data in conjunction with the discrete dynamic choice data to exploit cross-equation restrictions, whereas Taber does not.

The HN treatment of serially correlated unobservables is more general than any discussion that appears in the current dynamic discrete choice and dynamic treatment effect literature. They do not invoke the strong sequential conditional independence assumptions used in the dynamic treatment effect literature in statistics (Gill and Robins, 2001; Lechner and Miquel, 2002; Lok, 2007; Robins, 1989, 1997), nor do they invoke the closely related conditional temporal independence of unobserved state variables given observed state variables invoked by Rust (1987), Hotz and Miller (1988, 1993), Manski (1993) and Magnac and Thesmar (2002) (in the first part of their paper) or the independence assumptions invoked by Wolpin (1984).⁸⁴ HN allow for more general time series dependence in the

⁸³Pakes and Simpson (1989) sketch a proof of identification of a model of the option values of patents that is based on limit sets for an option model.

⁸⁴Manski (1993) and Hotz and Miller (1993) use a synthetic cohort effect approach that assumes that young agents will follow the transitions of contemporaneous older agents in making their lifecycle decisions. Manski and Hotz and Miller exclude any temporally dependent unobservables from their models. The synthetic cohort approach has been widely used in labor economics at least since Mincer (1974). See Ghez and Becker (1975), MaCurdy (1981) and

unobservables than is entertained by Pakes (1986), Keane and Wolpin (1997) or Eckstein and Wolpin (1999).⁸⁵

Like Miller (1984) and Pakes (1986), HN explicitly model, identify and estimate agent learning that affects expected future returns.⁸⁶ Pakes and Miller assume functional forms for the distributions of the error process and for the serial correlation pattern about information updating and time series dependence. The HN analysis of the unobservables is nonparametric and they estimate, rather than impose, the stochastic structure of the information updating process.

Virtually all papers in the literature, including the HN analysis, invoke rational expectations. An exception is the analysis of Manski (1993) who replaces rational expectations with a synthetic cohort assumption that choices and outcomes of one group can be observed (and acted on) by a younger group. This assumption is more plausible in stationary environments and excludes any temporal dependence in unobservables. In recent work, Manski (2004) advocates use of elicited expectations as an alternative to the synthetic cohort approach.

While HN use rational expectations, they estimate, rather than impose the structure of agent information sets. Miller (1984), Pakes (1986), Keane and Wolpin (1997), and Eckstein and Wolpin (1999) assume that they know the law governing the evolution of agent information up to unknown parameters.⁸⁷ Following the procedure presented in Cunha and Heckman (2007, 2008); Cunha, Heckman, and Navarro (2005, 2006) and Navarro (2005), HN can test for which factors (θ) appear in agent information sets at different stages of the life cycle and they identify the distributions of the unobservables nonparametrically.

The HN analysis of dynamic treatment effects is comparable, in some aspects, to the recent continuous-time event-history approach of Abbring and Van den Berg (2003b) previously analyzed. Those authors build a continuous-time model of counterfactuals for outcomes that are durations. They model treatment assignment times using a continuous-time duration model.

The HN analysis is in discrete time and builds on previous work by Heckman (1981a,c) on heterogeneity and state dependence that identifies the causal effect of employment (or unemployment) on future employment (or unemployment).⁸⁸ They model time to treatment and associated vectors of outcome equations that may be discrete, continuous or mixed discrete-continuous. In a discrete-time setting, they are able to generate a variety of distributions of counterfactuals and economically motivated parameters. They allow for heterogeneity in responses to treatment that has a general time series structure.

As noted in section 5.4, Abbring and Van den Berg (2003b) do not identify explicit agent

Mincer (1974) for applications of the synthetic cohort approach. For empirical evidence against the assumption that the earnings of older workers are a reliable guide to the earnings of younger workers in models of earnings and schooling choices for recent cohorts of workers, see Heckman, Lochner, and Todd (2006).

⁸⁵Rust (1994) provides a clear statement of the stochastic assumptions underlying the dynamic discrete-choice literature up to the date of his survey.

⁸⁶As previously noted, the previous literature assumes learning only about current costs.

⁸⁷They specify *a priori* particular processes of information arrival as well as which components of the unobservables agents know and act on, and which components they do not.

⁸⁸Heckman and Borjas (1980) investigate these issues in a continuous-time duration model. See also Heckman and MaCurdy (1980).

information sets as HN do in their paper and as is done in Cunha, Heckman, and Navarro (2005), and they do not model learning about future rewards. Their outcomes are restricted to be continuous-time durations. The HN framework is formulated in discrete time, which facilitates the specification of richer unobserved and observed covariate processes than those entertained in the continuous-time framework of Abbring and Van den Berg (2003b). It is straightforward to attach a vector of treatment outcomes in the HN model that includes continuous outcomes, discrete outcomes and durations expressed as binary strings.⁸⁹ At a practical level, the approach often can produce very fine-grained descriptions of continuous-time phenomena by using models with many finite periods. Clearly, a synthesis of the event-history approach with the HN approach would be highly desirable. That would entail taking continuous-time limits of the discrete-time models. It is a task that awaits completion.

Flinn and Heckman (1982) utilize information on stopping times and associated wages to derive cross-equation restrictions to partially identify an equilibrium job search model for a stationary economic environment where agents have an infinite horizon. They establish that the model is non-parametrically nonidentified. Their analysis shows that use of outcome data in conjunction with data on stopping times is not sufficient to secure nonparametric identification of a dynamic discrete-choice model, even when the reward function is linear in outcomes unlike the reward functions in Rust (1987) and Magnac and Thesmar (2002). Parametric restrictions can break their nonidentification result. Abbring and Campbell (2005) exploit such restrictions, together with cross-equation restrictions on stopping times and noisy outcome measures, to prove identification of an infinite-horizon model of firm survival and growth with entrepreneurial learning. Alternatively, nonstationarity arising from finite horizons can break their nonidentification result (see Wolpin, 1987). The HN analysis exploits the finite-horizon backward-induction structure of our model in conjunction with outcome data to secure identification and does not rely on arbitrary period by period exclusion restrictions. They substantially depart from the assumptions maintained in Rust’s nonidentification theorem (1994). They achieve identification by using cross-equation restrictions, linearity of preferences and additional measurements, and exploiting the structure of their finite horizon nonrecurrent model. Nonstationarity of regressors greatly facilitates identification by producing both exclusion and curvature restrictions which can substitute for standard exclusion restrictions.

6 Conclusion

This paper has surveyed recent approaches to using panel data to evaluate policies. We have compared and contrasted the statistical dynamic treatment approach based on sequential conditional independence assumptions that generalize matching to approaches developed in econometrics. We compared and contrasted a continuous-time event-history approach developed by Abbring and Van den Berg (2003b) to discrete-time reduced-form and structural models developed by Heckman and Navarro (2007), and Cunha, Heckman, and Navarro (2005).

⁸⁹Abbring (2008) considers nonparametric identification of mixed semi-Markov event-history models that extends his work with Van den Berg. See section 4.

References

- Aakvik, A., J. J. Heckman, and E. J. Vytlačil (2005). Estimating treatment effects for discrete outcomes when responses to treatment vary: An application to Norwegian vocational rehabilitation programs. *Journal of Econometrics* 125(1-2), 15–51.
- Aalen, O. O. and H. K. Gjessing (2004, December). Survival models based on the Ornstein-Uhlenbeck process. *Lifetime Data Analysis* 10(4), 407–423.
- Abbring, J. H. (2002, February). Stayers versus defecting movers: A note on the identification of defective duration models. *Economics Letters* 74(3), 327–331.
- Abbring, J. H. (2003). Dynamic econometric program evaluation. Discussion Paper 804, IZA, Bonn. Paper prepared for the H. Theil Memorial Conference, Amsterdam, 16-18 August 2002.
- Abbring, J. H. (2007). Mixed hitting-time models. Discussion Paper 07-57/3, Tinbergen Institute, Amsterdam.
- Abbring, J. H. (2008). The event-history approach to program evaluation. In D. Millimet, J. Smith, and E. Vytlačil (Eds.), *Modeling and Evaluating Treatment Effects in Econometrics*, Volume 21 of *Advances in Econometrics*. Oxford: Elsevier Science. Forthcoming.
- Abbring, J. H. and J. R. Campbell (2005). A firm’s first year. Discussion Paper 05-046/3, Tinbergen Institute, Amsterdam.
- Abbring, J. H. and J. J. Heckman (2007). Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6B, pp. 5145–5303. Amsterdam: Elsevier.
- Abbring, J. H. and G. J. Van den Berg (2003a, September). The identifiability of the mixed proportional hazards competing risks model. *Journal of the Royal Statistical Association, Series B* 65(3), 701–710.
- Abbring, J. H. and G. J. Van den Berg (2003b, September). The nonparametric identification of treatment effects in duration models. *Econometrica* 71(5), 1491–1517.
- Abbring, J. H. and G. J. Van den Berg (2003c). A simple procedure for inference on treatment effects in duration models. Discussion Paper 2003:19, IFAU, Uppsala.
- Abbring, J. H. and G. J. Van den Berg (2004, January). Analyzing the effect of dynamically assigned treatments using duration models, binary treatment models, and panel data models. *Empirical Economics* 29(1), 5–20.
- Abbring, J. H. and G. J. Van den Berg (2005). Social experiments and instrumental variables with duration outcomes. Discussion Paper 2005-047/3, Tinbergen Institute, Amsterdam.

- Abbring, J. H., G. J. Van den Berg, and J. C. Van Ours (2005, July). The effect of unemployment insurance sanctions on the transition rate from unemployment to employment. *Economic Journal* 115(505), 602–630.
- Aldrich, J. (1989, January). Autonomy. *Oxford Economic Papers* 41(1), 15–34.
- Andersen, P. K., Ø. Borgan, R. Gill, and N. Keiding (1993). *Statistical Models Based on Counting Processes*. New York: Springer-Verlag.
- Belzil, C. and J. Hansen (2002, September). Unobserved ability and the return to schooling. *Econometrica* 70(5), 2075–2091.
- Black, D. A., J. A. Smith, M. C. Berger, and B. J. Noel (2003, September). Is the threat of reemployment services more effective than the services themselves? Evidence from random assignment in the UI system. *American Economic Review* 93(4), 1313–1327.
- Bonnal, L., D. Fougère, and A. Sérandon (1997, October). Evaluating the impact of French employment policies on individual labour market histories. *Review of Economic Studies* 64(4), 683–713.
- Cameron, S. V. and J. J. Heckman (1998, April). Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of American males. *Journal of Political Economy* 106(2), 262–333.
- Card, D. and D. G. Sullivan (1988, May). Measuring the effect of subsidized training programs on movements in and out of employment. *Econometrica* 56(3), 497–530.
- Carneiro, P., K. Hansen, and J. J. Heckman (2001, Fall). Removing the veil of ignorance in assessing the distributional impacts of social policies. *Swedish Economic Policy Review* 8(2), 273–301.
- Carneiro, P., K. Hansen, and J. J. Heckman (2003, May). Estimating distributions of treatment effects with an application to the returns to schooling and measurement of the effects of uncertainty on college choice. *International Economic Review* 44(2), 361–422.
- Chan, T. Y. and B. H. Hamilton (2006, December). Learning, private information and the economic evaluation of randomized experiments. *Journal of Political Economy* 114(6), 997–1040.
- Cunha, F. and J. J. Heckman (2007). The evolution of uncertainty in labor earnings in the U.S. economy. Unpublished manuscript, University of Chicago. Under revision.
- Cunha, F. and J. J. Heckman (2008). A new framework for the analysis of inequality. *Macroeconomic Dynamics*, forthcoming.
- Cunha, F., J. J. Heckman, and S. Navarro (2005, April). Separating uncertainty from heterogeneity in life cycle earnings, The 2004 Hicks Lecture. *Oxford Economic Papers* 57(2), 191–261.

- Cunha, F., J. J. Heckman, and S. Navarro (2006). Counterfactual analysis of inequality and social mobility. In S. L. Morgan, D. B. Grusky, and G. S. Fields (Eds.), *Mobility and Inequality: Frontiers of Research in Sociology and Economics*, Chapter 4, pp. 290–348. Stanford, CA: Stanford University Press.
- Cunha, F., J. J. Heckman, and S. Navarro (2007, November). The identification and economic content of ordered choice models with stochastic cutoffs. *International Economic Review* 48(4), 1273–1309.
- Eberwein, C., J. C. Ham, and R. J. LaLonde (1997, October). The impact of being offered and receiving classroom training on the employment histories of disadvantaged women: Evidence from experimental data. *Review of Economic Studies* 64(4), 655–682.
- Eckstein, Z. and K. I. Wolpin (1999, November). Why youths drop out of high school: The impact of preferences, opportunities, and abilities. *Econometrica* 67(6), 1295–1339.
- Falmagne, J.-C. (1985). *Elements of Psychophysical Theory*. Oxford Psychology Series No. 6. New York: Oxford University Press.
- Fitzenberger, B., A. Osikominu, and R. Völter (2006, May). Get training or wait? Long-run employment effects of training programs for the unemployed in West Germany. Discussion paper 2121, IZA, Bonn.
- Fleming, T. R. and D. P. Harrington (1991). *Counting Processes and Survival Analysis*. New York: Wiley.
- Flinn, C. and J. J. Heckman (1982, January). New methods for analyzing structural models of labor force dynamics. *Journal of Econometrics* 18(1), 115–68.
- Florens, J. and M. Mouchart (1982). A note on noncausality. *Econometrica* 50, 583–591.
- Freedman, D. A. (2004, August). On specifying graphical models for causation and the identification problem. *Evaluation Review* 28(4), 267–293.
- Freund, J. E. (1961, December). A bivariate extension of the exponential distribution. *Journal of the American Statistical Association* 56(296), 971–977.
- Frisch, R. (1938). Autonomy of economic relations. Paper given at League of Nations. Reprinted in D.F. Hendry and M.S. Morgan (1995), *The Foundations of Econometric Analysis*, Cambridge University Press.
- Ghez, G. R. and G. S. Becker (1975). *The Allocation of Time and Goods over the Life Cycle*. New York: National Bureau of Economic Research.
- Gill, R. D. and J. M. Robins (2001, December). Causal inference for complex longitudinal data: The continuous case. *Annals of Statistics* 29(6), 1785–1811.

- Granger, C. W. J. (1969, August). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37(3), 424–438.
- Gritz, R. M. (1993, May-June). The impact of training on the frequency and duration of employment. *Journal of Econometrics* 57(1-3), 21–51.
- Grubb, D. (2000). Eligibility criteria for unemployment benefits. In *OECD Economic Studies, Special issue: Making Work Pay*, Number 31, pp. 147–184. OECD.
- Haavelmo, T. (1943, January). The statistical implications of a system of simultaneous equations. *Econometrica* 11(1), 1–12.
- Ham, J. C. and R. J. LaLonde (1996, January). The effect of sample selection and initial conditions in duration models: Evidence from experimental data on training. *Econometrica* 64(1), 175–205.
- Hansen, L. P. and T. J. Sargent (1980, February). Formulating and estimating dynamic linear rational expectations models. *Journal of Economic Dynamics and Control* 2(1), 7–46.
- Heckman, J. J. (1974, July). Shadow prices, market wages, and labor supply. *Econometrica* 42(4), 679–694.
- Heckman, J. J. (1981a). Heterogeneity and state dependence. In S. Rosen (Ed.), *Studies in Labor Markets*, National Bureau of Economic Research, pp. 91–139. University of Chicago Press.
- Heckman, J. J. (1981b). The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process and some Monte Carlo evidence. In C. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*, pp. 179–85. Cambridge, MA: MIT Press.
- Heckman, J. J. (1981c). Statistical models for discrete panel data. In C. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*, pp. 114–178. Cambridge, MA: MIT Press.
- Heckman, J. J. (1990, May). Varieties of selection bias. *American Economic Review* 80(2), 313–318.
- Heckman, J. J. (2005, August). The scientific model of causality. *Sociological Methodology* 35(1), 1–97.
- Heckman, J. J. and G. J. Borjas (1980, August). Does unemployment cause future unemployment? Definitions, questions and answers from a continuous time model of heterogeneity and state dependence. *Economica* 47(187), 247–283. Special Issue on Unemployment.
- Heckman, J. J. and B. E. Honoré (1989, June). The identifiability of the competing risks model. *Biometrika* 76(2), 325–330.
- Heckman, J. J. and B. E. Honoré (1990, September). The empirical content of the Roy model. *Econometrica* 58(5), 1121–1149.

- Heckman, J. J., R. J. LaLonde, and J. A. Smith (1999). The economics and econometrics of active labor market programs. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 3A, Chapter 31, pp. 1865–2097. New York: North-Holland.
- Heckman, J. J., L. J. Lochner, and C. Taber (1998, January). Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Review of Economic Dynamics* 1(1), 1–58.
- Heckman, J. J., L. J. Lochner, and P. E. Todd (2006). Earnings equations and rates of return: The Mincer equation and beyond. In E. A. Hanushek and F. Welch (Eds.), *Handbook of the Economics of Education*, Chapter 7, pp. 307–458. Amsterdam: Elsevier.
- Heckman, J. J. and T. E. MaCurdy (1980, January). A life cycle model of female labour supply. *Review of Economic Studies* 47(1), 47–74.
- Heckman, J. J. and S. Navarro (2004, February). Using matching, instrumental variables, and control functions to estimate economic choice models. *Review of Economics and Statistics* 86(1), 30–57.
- Heckman, J. J. and S. Navarro (2005). Empirical estimates of option values of education and information sets in a dynamic sequential choice model. Unpublished manuscript, University of Chicago, Department of Economics.
- Heckman, J. J. and S. Navarro (2007, February). Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics* 136(2), 341–396.
- Heckman, J. J. and B. S. Singer (1984, January-February). Econometric duration analysis. *Journal of Econometrics* 24(1-2), 63–132.
- Heckman, J. J. and B. S. Singer (1986). Econometric analysis of longitudinal data. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics*, Volume 3, Chapter 29, pp. 1690–1763. North-Holland.
- Heckman, J. J. and J. A. Smith (1998). Evaluating the welfare state. In S. Strom (Ed.), *Econometrics and Economic Theory in the Twentieth Century: The Ragnar Frisch Centennial Symposium*, pp. 241–318. New York: Cambridge University Press.
- Heckman, J. J., J. Stixrud, and S. Urzua (2006, July). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics* 24(3), 411–482.
- Heckman, J. J. and C. Taber (1994). Econometric mixture models and more general models for unobservables in duration analysis. *Statistical Methods in Medical Research* 3(3), 279–299.
- Heckman, J. J., J. L. Tobias, and E. J. Vytlačil (2001, October). Four parameters of interest in the evaluation of social programs. *Southern Economic Journal* 68(2), 210–223.

- Heckman, J. J., J. L. Tobias, and E. J. Vytlacil (2003, August). Simple estimators for treatment parameters in a latent variable framework. *Review of Economics and Statistics* 85(3), 748–754.
- Heckman, J. J., S. Urzua, and E. J. Vytlacil (2006). Understanding instrumental variables in models with essential heterogeneity. *Review of Economics and Statistics* 88(3), 389–432.
- Heckman, J. J., S. Urzua, and G. Yates (2007). The identification and estimation of option values in a model with recurrent states. Unpublished manuscript, University of Chicago, Department of Economics.
- Heckman, J. J. and E. J. Vytlacil (1999, April). Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences* 96, 4730–4734.
- Heckman, J. J. and E. J. Vytlacil (2001). Causal parameters, treatment effects and randomization. Unpublished manuscript, University of Chicago, Department of Economics.
- Heckman, J. J. and E. J. Vytlacil (2005, May). Structural equations, treatment effects and econometric policy evaluation. *Econometrica* 73(3), 669–738.
- Heckman, J. J. and E. J. Vytlacil (2007a). Econometric evaluation of social programs, part I: Causal models, structural models and econometric policy evaluation. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6B, pp. 4779–4874. Amsterdam: Elsevier.
- Heckman, J. J. and E. J. Vytlacil (2007b). Econometric evaluation of social programs, part II: Using the marginal treatment effect to organize alternative economic estimators to evaluate social programs and to forecast their effects in new environments. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6B, pp. 4875–5144. Amsterdam: Elsevier.
- Hendry, D. F. and M. S. Morgan (1995). *The Foundations of Econometric Analysis*. New York: Cambridge University Press.
- Holland, P. W. (1986, December). Statistics and causal inference. *Journal of the American Statistical Association* 81(396), 945–960.
- Honoré, B. E. (1993, January). Identification results for duration models with multiple spells. *Review of Economic Studies* 60(1), 241–246.
- Honoré, B. E. and A. Lewbel (2002, September). Semiparametric binary choice panel data models without strictly exogenous regressors. *Econometrica* 70(5), 2053–2063.
- Hotz, V. J. and R. A. Miller (1988, January). An empirical analysis of life cycle fertility and female labor supply. *Econometrica* 56(1), 91–118.
- Hotz, V. J. and R. A. Miller (1993, July). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60(3), 497–529.

- Hurwicz, L. (1962). On the structural form of interdependent systems. In E. Nagel, P. Suppes, and A. Tarski (Eds.), *Logic, Methodology and Philosophy of Science*, pp. 232–239. Stanford University Press.
- Kalbfleisch, J. D. and R. L. Prentice (1980). *The Statistical Analysis of Failure Time Data*. New York: Wiley.
- Keane, M. P. and K. I. Wolpin (1994, November). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte Carlo evidence. *The Review of Economics and Statistics* 76(4), 648–672.
- Keane, M. P. and K. I. Wolpin (1997, June). The career decisions of young men. *Journal of Political Economy* 105(3), 473–522.
- Keiding, N. (1999, September). Event history analysis and inference from observational epidemiology. *Statistics in Medicine* 18(17-18), 2353–2363.
- Lancaster, T. (1979, July). Econometric methods for the duration of unemployment. *Econometrica* 47(4), 939–956.
- Leamer, E. E. (1985, Spring). Vector autoregressions for causal inference? *Carnegie-Rochester Conference Series on Public Policy* 22, 255–303.
- Lechner, M. and R. Miquel (2002). Identification of effects of dynamic treatments by sequential conditional independence assumptions. Discussion paper, University of St. Gallen, Department of Economics.
- Lewbel, A. (2000, July). Semiparametric qualitative response model estimation with unknown heteroscedasticity or instrumental variables. *Journal of Econometrics* 97(1), 145–177.
- Lok, J. J. (2007). Statistical modelling of causal effects in time. *Annals of Statistics*. In press.
- MaCurdy, T. E. (1981, December). An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy* 89(6), 1059–1085.
- Magnac, T. and D. Thesmar (2002, March). Identifying dynamic discrete decision processes. *Econometrica* 70(2), 801–816.
- Manski, C. F. (1988, September). Identification of binary response models. *Journal of the American Statistical Association* 83(403), 729–738.
- Manski, C. F. (1993, July). Dynamic choice in social settings: Learning from the experiences of others. *Journal of Econometrics* 58(1-2), 121–136.
- Manski, C. F. (2004, September). Measuring expectations. *Econometrica* 72(5), 1329–1376.

- Matzkin, R. L. (1992, March). Nonparametric and distribution-free estimation of the binary threshold crossing and the binary choice models. *Econometrica* 60(2), 239–270.
- Matzkin, R. L. (1993, July). Nonparametric identification and estimation of polychotomous choice models. *Journal of Econometrics* 58(1-2), 137–168.
- Matzkin, R. L. (1994). Restrictions of economic theory in nonparametric methods. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, pp. 2523–58. New York: North-Holland.
- Matzkin, R. L. (2003, September). Nonparametric estimation of nonadditive random functions. *Econometrica* 71(5), 1339–1375.
- Meyer, B. D. (1990, July). Unemployment insurance and unemployment spells. *Econometrica* 58(4), 757–782.
- Meyer, B. D. (1996, January). What have we learned from the Illinois reemployment bonus experiment? *Journal of Labor Economics* 14(1), 26–51.
- Miller, R. A. (1984, December). Job matching and occupational choice. *Journal of Political Economy* 92(6), 1086–1120.
- Mincer, J. (1974). *Schooling, Experience and Earnings*. New York: Columbia University Press for National Bureau of Economic Research.
- Mortensen, D. T. (1977, July). Unemployment insurance and job search decisions. *Industrial and Labor Relations Review* 30(4), 505–517.
- Mortensen, D. T. (1986). Job search and labor market analysis. In O. Ashenfelter and P. R. C. Layard (Eds.), *Handbook of Labor Economics*, Volume 2 of *Handbooks in Economics*, pp. 849–919. New York: Elsevier Science.
- Murphy, S. A. (2003, May). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Association, Series B* 65(2), 331–366.
- Navarro, S. (2005). *Understanding Schooling: Using Observed Choices to Infer Agent’s Information in a Dynamic Model of Schooling Choice When Consumption Allocation is Subject to Borrowing Constraints*. Ph.D. Dissertation, University of Chicago, Chicago, IL.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. essay on principles. *Roczniki Nauk Rolniczych* 10, 1–51. in Polish; edited and translated version of Section 9 by D.M. Dabrowska and T.P. Speed (1990), *Statistical Science*, 5, 465–472.
- Pakes, A. (1986, July). Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica* 54(4), 755–784.

- Pakes, A. and M. Simpson (1989). Patent renewal data. *Brookings Papers on Economic Activity* (Special Issue), 331–401.
- Pearl, J. (2000). *Causality*. Cambridge, England: Cambridge University Press.
- Quandt, R. E. (1958, December). The estimation of the parameters of a linear regression system obeying two separate regimes. *Journal of the American Statistical Association* 53(284), 873–880.
- Quandt, R. E. (1972, June). A new approach to estimating switching regressions. *Journal of the American Statistical Association* 67(338), 306–310.
- Ridder, G. (1986, April). An event history approach to the evaluation of training, recruitment and employment programmes. *Journal of Applied Econometrics* 1(2), 109–126.
- Ridder, G. (1990, April). The non-parametric identification of generalized accelerated failure-time models. *Review of Economic Studies* 57(2), 167–181.
- Robins, J. M. (1989). The analysis of randomized and non-randomized AIDS treatment trials using a new approach to causal inference in longitudinal studies. In L. Sechrest, H. Freeman, and A. Mulley (Eds.), *Health Services Research Methodology: A Focus on AIDS*, pp. 113–159. Rockville, MD: United States Department of Health and Human Services, National Center for Health Services Research and Health Care Technology Assessment.
- Robins, J. M. (1997). Causal inference from complex longitudinal data. In M. Berkane (Ed.), *Latent Variable Modeling and Applications to Causality. Lecture Notes in Statistics*, pp. 69–117. New York: Springer-Verlag.
- Rosenzweig, M. R. and K. I. Wolpin (2000, December). Natural ”natural experiments” in economics. *Journal of Economic Literature* 38(4), 827–874.
- Roy, A. (1951, June). Some thoughts on the distribution of earnings. *Oxford Economic Papers* 3(2), 135–146.
- Rubin, D. B. (1974, October). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66(5), 688–701.
- Rubin, D. B. (1986). Statistics and causal inference: Comment: Which ifs have causal answers. *Journal of the American Statistical Association* 81(396), 961–962.
- Rust, J. (1987, September). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55(5), 999–1033.
- Rust, J. (1994). Structural estimation of Markov decision processes. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics, Volume*, pp. 3081–3143. New York: North-Holland.
- Taber, C. R. (2000, June). Semiparametric identification and heterogeneity in discrete choice dynamic programming models. *Journal of Econometrics* 96(2), 201–229.

- Thurstone, L. L. (1930). *The Fundamentals of Statistics*. New York: The Macmillan Company.
- Van den Berg, G. J. (2001). Duration models: Specification, identification and multiple durations. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics*, Volume 5 of *Handbooks in Economics*, pp. 3381–3460. New York: North-Holland.
- Van den Berg, G. J., A. Holm, and J. C. Van Ours (2002, November). Do stepping-stone jobs exist? Early career paths in the medical profession. *Journal of Population Economics* 15(4), 647–665.
- Van den Berg, G. J., B. Van der Klaauw, and J. C. Van Ours (2004, January). Punitive sanctions and the transition rate from welfare to work. *Journal of Labor Economics* 22(1), 211–241.
- Van der Laan, M. J. and J. M. Robins (2003). *Unified Methods for Censored Longitudinal Data and Causality*. New York: Springer-Verlag.
- Wolpin, K. I. (1984, October). An estimable dynamic stochastic model of fertility and child mortality. *Journal of Political Economy* 92(5), 852–874.
- Wolpin, K. I. (1987, July). Estimating a structural search model: The transition from school to work. *Econometrica* 55(4), 801–817.

List of figures

1. Sample distribution of schooling attainment probabilities for males from the National Longitudinal Survey of Youth. (Source: Heckman and Navarro (2007))