# Dynamic Social Experiments

Christopher Ferrall

Department of Economics
Queen's University

`http://econ.queensu.ca/ ferrall`

January 6, 2006

**Abstract:**

This paper develops a framework for analyzing the outcome of experiments carried out on forward-looking subjects. Natural experiments, unexpected policy changes, and true experiments are all included in the framework as special cases. These concepts are defined in conjunction with explicit notions of controlled and randomized experiments. The persistent issues of sample-selection bias and heterogeneous impacts that surround interpretations of experiments are endogenous to the model. Special attention is given to interpreting empirical impact of the treatment within the model. The environments in which estimated mean impacts correspond to mean subjective impacts are specified, and they are found to be a small, uninteresting subset of environments contained within the framework.

**JEL Classification:** C1, C9

**Keywords:**
Dynamic Programming, Treatment Effects, Policy Experiments

**Correspondence:**

C. Ferrall, Dept. of Economics, Queen's University, Kingston, Ontario K7L 3N6, CANADA.

ferrallc@post.queensu.ca

# I. Introduction

The Negative Income Tax (NIT) experiments of the 1960s and 1970s were expected to set economic policy analysis on a course of ever increasing confidence and precision (Orcutt and Orcutt 1968). Yet, based on the methods and resources available at the time, the NIT experiments failed to provide definitive measurement of the elasticity of labor supply (Pencavel 1987). The experimental design did not clearly identify the parameters of interest, and a host of econometric problems were encountered. This experience and the steady growth of non-experimental data sources and econometric methods kept large-scale experiments from becoming commonplace. More recently economists have become less ambivalent about experiments, in part because the goals stated for social experiments have become more modest. Social experiments are now more carefully designed to demonstrate the impact of a particular policy change on selected outcomes. Indeed, despite the failures of the NIT experiments, the potential power of controlled experiments to solve problems of endogenous explanatory variables has influenced non-experimental econometric analysis by placing great emphasis on so-called natural experiments and the closely linked technique of instrumental variables.

This paper revisits the use of social experiments to identify parameters of a model of individual behavior.[1] In contrast to the state-of-the-art at the time of the NIT, current empirical models of behavior emphasize forward-looking decision-making, unobserved differences between agents, and differences between the information available to agents and the econometrician. These considerations lead to modeling the behavior of subjects in an experiments as a discrete-choice dynamic program. Thus this paper relates closely to Rust (1994) which characterized the estimation of discrete-choice Markov decision problems. Here three state variables are introduced to Rust's framework to account for

---

[1] The leading example of a controlled experiment designed to study and to be studied by a unified model of individual behavior is Shearer's (1994) experiment on incentives to carry out hard physical labor (tree planting).

experimental (or quasi-experimental) treatment. The standard model is also extended to integrate unanticipated ("zero-probability") events, sample selection, observed and unobserved heterogeneity, averaged data, and hypothetical uses of the results. While nearly all these elements have appeared before in applications of discrete choice dynamic programming, they have been seen as extensions to the standard model of a homogeneous, stable environment with random sampling and fixed initial conditions. Here the elements of heterogeneity, exogenous variation, and endogenous initial conditions are combined within a unified strategy to solve, estimate, and apply a model.

Much of the recent literature on experiments and program evaluation seeks to define quantities, such as the mean treatment effect, that can be recovered from data under weak assumptions and a trend toward viewing such quantities as primitive parameters. For example, in discussing the "local average treatment effect" (LATE) introduced by Imbens and Angrist (1994), Manksi (1996) notes, "Imbens and Angrist do not argue that an analyst should…be concerned with this subpopulation. They motivate their interest [in LATE] by stressing that this quantity is identified in circumstances where the effect of treatment on compliers is not identified" (p. 723). In contrast, this paper reserves the term 'parameter' for unknown constants of two types: a parameter of an agent's objective function or a proportion of a (unknown) agent type in the population. In short, the objects of interest are independent of the nature of the experiment. Quantities such as mean treatment effects, which are dependent of the experiment, are derived from the underlying parameters and the policy environment surrounding agents.

Several implications are presented suggesting that intuition from a static model and a view of individual treatment impacts as primitive parameters does not readily extend into the dynamic world in which individual impact arises from individual response. In short, while random assignment to treatment alone is useful, it is not sufficient on its own to identify elements of the environment that are important for policy analysis. This paper is not the first to suggest that experimental impacts, while easy to interpret, may be of limited

2

value unless combined with a model that accounts for how subjects become eligible for random assignment and why they respond to treatment. Others include Heckman et al (1999), Hotz et. al (1999), Ichimura and Taber (2000), Rosenzweig and Wolpin (2000), Todd and Wolpin (2003) and Lise et al (2004) and Moffitt (2004).

The framework presented here encompasses provides a formal language for describing experiments and using all the information available in the outcomes they generate. The formal definition of a social experiment is sufficiently precise to describe a complete solution algorithm, alternative estimation procedures, and use of results to analyze hypothetical policy changes. The definition is also sufficiently general to include designed experiments, field trials, demonstration projects, unexpected policy changes, and random acts of nature. As examples, four situations are mapped into the basic framework to illustrate its ability to represent exogenous variation as it is usually portrayed in the literature.

## II. The Social Environment

Before describing an experiment it is necessary to describe the environment in which the experiment takes place. The environment is called 'social' not because individuals interact, although such interactions can be incorporated into the framework (see VII). Rather, policies parameters and other aspects of the setting outside the experiment affect the results generated by the experiment, and in this sense the social situation affects the outcome of an experiment. (The Appendix provides some notation and convention used in the technical aspects of the paper.)

### II.A Subject Behavior

The model of subject behavior is a discrete-state, discrete-choice dynamic program.[2]

---

[2] The framework can be extended to include approximation of a state space where some state variables take on a continuum of values using methods described in, for example,

Each period the subject is in a state $\theta$, an element of the state space $\Theta$. Each period the subject chooses an action vector $\alpha$ from a set of feasible actions $\mathbf{A}(\theta)$. A choice-state combination $(\alpha, \theta)$ is referred to as an *outcome*. The value of an outcome is written

$$v(\alpha, \theta) = U(\alpha, \theta) + \delta E\left[V(\theta')\right] = U(\alpha, \theta) + \delta \sum_{\theta'} P\{\theta' \,|\, \alpha\,,\, \theta\} V(\theta'), \tag{1}$$

and depends on the one-period utility $U(\alpha, \theta)$, the discount factor $\delta$, and the outcome-to-state transition probability $P\{\theta' \,|\, \alpha\,,\, \theta\}$. $V(\theta')$ is the value of entering state $\theta'$ given that an optimal decision will be made then and for all future periods:

$$\forall \theta \in \Theta, \qquad V(\theta) = max_{\alpha \in \mathbf{A}(\theta)} \quad v(\alpha, \theta). \tag{2}$$

The dynamic program in (1) and (2) can describe environments in which no unexpected events occur to subjects. Yet the essence of an experiment is that subjects do not anticipate treatment. More notation is required to account for the exogenous and unexpected aspects of an experiment. For this purpose it proves useful to think of the state vector $\theta$ as the concatenation of five sub-vectors, each of which contains state variables that play specific roles in the model:

$$\theta \equiv \begin{pmatrix} \theta_{\text{clock}} & \theta_{\text{exp}} & \theta_{\text{end}} & \theta_{\text{exog}} & \theta_{\text{pol}} \end{pmatrix} \tag{3}$$

$$= \Big( \begin{pmatrix} t & r & f \end{pmatrix} \quad \begin{pmatrix} g & e \end{pmatrix} \quad \begin{pmatrix} \cdots & d & k \end{pmatrix} \quad \begin{pmatrix} \Lambda & \Gamma \end{pmatrix} \quad \begin{bmatrix} \Psi_p[d] \end{bmatrix} \Big).$$

The 'pol' vector contains parameters describing current public policies that affect subjects. The 'exog' vector contains preference and technology parameters; 'end' contains the variables typically thought of as state variables, i.e. those endogenous to the subject's problem (1)-(2). These variables pertain to the world outside the experiment. The ellipsis indicates that the modeler chooses the elements of the endogenous vector except for the required variables $d$ and $k$, which are described below. The remaining two sub-vectors

Rust (1997). A continuous state space is often assumed so as to smooth the model's predictions. Other approaches to smoothness are integrated into this framework and may make a continuous state space unnecessary.

describe the subject's situation inside the experiment, and they are defined in the next section. Outside the experiment the values of $\theta_{\text{clock}}$ and $\theta_{\text{exp}}$ are transparent to the subject. The required state variables within each sub-vector that are listed above will be defined as the role of each sub-vector is described.

When the state and action spaces are large, the transition $P\{\theta' \mid \alpha, \theta\}$ becomes exceedingly complex if allowed to take an arbitrary form. On the other hand, any restrictions on the transition must not rule out common specifications. To balance these concerns the evolution of state variables is described with some special notation.

**Definition 1.** A **discrete jump process** $q' = q^{\star}(\bar{q}, [\pi_j], [\mathbf{Q}_j])$ means $\forall\, (\alpha, \theta)$,

$$P\{q' \mid \alpha, \theta\} = \sum_{j=1}^{J(\alpha,\theta)} \pi_j(\alpha, \theta) \frac{\mathcal{B}\left[q' \in \mathbf{Q}_j(\alpha, \theta)\right]}{|\mathbf{Q}_j(\alpha, \theta)|} + \left(1 - \sum_{j}^{J(\alpha,\theta)} \pi_j(\alpha, \theta)\right) \mathcal{B}\left[q' = \bar{q}(\alpha, \theta)\right]. \tag{4}$$

In words, $q^{\star}(\bar{q}, [\pi_j], [\mathbf{Q}_j])$ means that the next realized value of $q$ is either the default value $\bar{q}$ (a scalar) or jumps into one of $J$ sets $\mathbf{Q}_j$ with corresponding probability $\pi_j$. The notation $[\mathbf{Q}_j]$ emphasizes that the jump sets are organized into a vector parallel to the vector of probabilities, $[\pi_j]$. The probability of the default event is $1 - \sum_j \pi_j$, although the default value can appear in the jump sets, so $\bar{q}$ may be more likely to occur than the default event itself. Within each jump set the values are equally likely: $\pi_j/|\mathbf{Q}_j|$. This is assumed without loss of generality, because the modeler has the option of specifying each jump set as a singleton. Note that (4) makes it explicit that the arguments of the jump process depend on the current outcome $(\alpha, \theta)$. In the extreme, each set could contain a single value, although most models have much more restrictive processes and it would be more difficult to classify different types of state variables without notation that succinctly describes the simple cases that occur most often, as this list of examples illustrates:[3]

---

[3] To show the flexibility of this notation, consider one extended example. Suppose that $q$ is iid at states where some other variable $z$ is 0. Otherwise, $q$ retains its current value. For example, $q$ might be an index for job offers in a search model. While unemployed ($z = 0$), job offers arrive iid from a discrete distribution, but while employed ($z = 1$) the offer stays constant at its current value. Using the discrete jump notation, this process is described

**Examples** A discrete jump process $q \in \theta$ is:

⋄ *absorbing at $n$* if $q = n \ \rightarrow \ q' = q^\star(n, [0], [\emptyset])$;

⋄ *invariant* if $q$ is absorbing for all values $q = 1, 2, \ldots, Q$;

⋄ *autonomous [at $\theta$]* if [at $\theta$] $q^\star$ is not a function of $\alpha$;

⋄ *independent [at $q = n$]* if ($q = n$ implies) $q^\star$ is not a function of $\alpha$ or any other state variable $s \neq q$;

⋄ *ergodic [at $\theta$]* if for $n = 1, 2, \ldots, Q$, $\exists\, \alpha^n \in \mathbf{A}(\theta) : P\{q' = n \mid (\alpha^n, \theta)\} > 0$;

⋄ *iid* if for constants $[p_i]$, $q' = q^\star\left( q, \begin{bmatrix} p_1 & \cdots & 1 - \sum_{i=1}^{Q-1} p_i \end{bmatrix}, \begin{bmatrix} \{1\} & \cdots & \{Q\} \end{bmatrix} \right)$;

⋄ *dependent* if $q$ is not ergodic, but for any two values $q^\star, q^{\star\star}$, there exists a chain of outcomes $(\alpha^1, \theta^1), (\alpha^2, \theta^2), \ldots, (\alpha^m, \theta^m)$, $m < Q - 1$, with corresponding values of $q$ such that $P\left\{q' = q^{\star\star} \mid (\alpha^m, \theta^m)\right\} > 0$, $P\left\{q' = q_m \mid (\alpha^{m-1}, \theta^{m-1})\right\} > 0$, $\ldots$, $P\left\{q' = q_1 \mid (\alpha^\star, \theta^\star)\right\} > 0$;

Invariant states as defined above differ across individuals but are fixed for a given individual. The set of invariant states partitions the state space. That is, there exist a set of subsets in which all the invariant states are constant, and the state of any subject in one of these subsets remains in the subset forever.

### Requirement 1. Transitions and Choice Probabilities.

**R1a.** $0 < \delta < 1$; $U(\alpha, \theta)$ is a smooth function of a vector of exogenous parameters $\Upsilon \in \mathbf{U} \subset \Re^{N1}$; $P\{\theta' \mid \alpha, \theta\}$ is a smooth function of a vector $\Pi \in \mathbf{P} \subset \Re^{N2}$ of exogenous parameters; $\mathbf{U}$ and $\mathbf{P}$ are bounded open sets.

---

as

$$q^\star\left( q, \mathcal{B}\,[z = 0] \begin{bmatrix} p_1 & p_2 & \cdots & 1 - \sum_i^{Q-1} p_i \end{bmatrix}, \begin{bmatrix} \{1\} & \{2\} & \cdots & \{Q\} \end{bmatrix} \right). \qquad (5)$$

Multiplying the vector of jump probabilities by the scalar $\mathcal{B}\,[z = 0]$ makes the probability of the default (and current) value equal 1 when $z \neq 0$. If the model specifies that $z$ is ergodic, this means that for any employed state there exists at least one action that will lead to a positive probability of unemployment. Here $q$ is dependent on $z$ at employed states. Although $q$ itself is not ergodic, its dependence on $z$ keeps the overall transition ergodic.

**R1b.** For all $\Pi \in \mathbf{P}$ each variable $q \in \theta$ follows a discrete jump process that is invariant, ergodic, or dependent.

**R1c.** State-contingent choice probabilities are smoothed by a logistic kernel with parameter $0 < \rho < 1$:

$$\tilde{v}(\alpha,\theta) \equiv \mathcal{B}\left[\alpha \in \mathbf{A}(\theta)\right] \exp\left\{\frac{\rho}{1-\rho}\left[v(\alpha,\theta) - V(\theta)\right]\right\}$$

$$P\{\alpha \mid \theta\} = \frac{\tilde{v}(\alpha,\theta)}{\sum_{\alpha'} \tilde{v}(\alpha',\theta)}. \tag{6}$$

From R1 the complete transition function can be described. The formulas for $P\{\theta' \mid \alpha, \theta\}$ and $P_s\{\theta' \mid \theta\}$ are given in the Appendix.

## II.B Observed and Unobserved Heterogeneity

A group of subjects that, for the purposes of the model and the experiment, share observable characteristics that do not vary over time is called a demographic group and is indexed by $d \in \theta_{\text{end}}$. The number of demographic groups, $D$, is pre-determined and finite. In addition, subjects are also divided into $K$ pre-determined groups indexed by $k \in \theta_{\text{end}}$ and sharing unobserved characteristics. A subject's demographic group and unobserved type completely specify their environment outside the experiment. For example, suppose a model is posited in which income tax rates depend on location but a subject's location is not determined endogenously by $\alpha$ and is taken as given in the analysis. Location is then a demographic variable. Each unique combination of location and the values of other demographic variables would define a demographic group. Further, the model posits that subjects differ in the marginal utility of leisure, which is also treated as given in the model but not directly observable to others. Neither government policy nor the experiment cannot directly treat subjects differently based on marginal utility, so each value of it would help define an unobserved type.

Individuals who differ in observed ways are likely to differ in unobserved ways. Although $d$ is an index of fixed exogenous variation, it is problematic to limit the correlation

7

between it and the underlying preference parameters indexed by $k$. Continuing the marginal tax rate example, people with low marginal utility of leisure may choose to live in locales with low income taxes. One solution is to model all results of interest including $d$ itself. This may not be feasible, and it can detract from a focus on the experimental results. A second solution is to let exogenous variables be parametric functions of observed characteristics and other 'deeper' estimated parameters. This creates a hybrid model in which some parameters are related to observed characteristics and some are not. The middle-ground strategy adopted here assumes non-parametric unobserved heterogeneity (Heckman and Singer 1984) and mixture probabilities that depend on the demographic index $d$. The advantages of this assumption over a hybrid model will be described below.

### Requirement 2. Exogenous Parameters.

**R2a.** $\theta_{\mathrm{end}}$ contains *exactly* two invariant states, denoted $d$ and $k$, and taking on positive integer values up to $D$ and $K$, respectively.

**R2b.** The exogenous parameters of a subject's problem (1)-(2) are collected in a vector of size $N = N_1 + N_2 + 2$ that is specific to each type $k$:

$$\Gamma[k] \equiv \begin{pmatrix} \Upsilon_k & \Pi_k & \delta_k & \rho_k \end{pmatrix}. \tag{7}$$

**R2c.** $U(\alpha, \theta)$ and $P\{\theta' \,|\, \alpha, \theta\}$ are (not necessarily smooth) functions of a vector $\Psi_p$ of policy instruments. The policy vector $\theta_{\mathrm{pol}}$ takes on exactly one value (and so is invariant) and consists of $D$ sub-vectors indexed by $d$:

$$\theta_{\mathrm{pol}} \equiv \begin{pmatrix} \Psi_p[1] & \Psi_p[2] & \cdots & \Psi_p[D] \end{pmatrix}. \tag{8}$$

**R2d.** The (invariant) proportion of the population with a given combination $\begin{pmatrix} k & d \end{pmatrix}$ equals $\bar{\lambda}[d]\lambda[d,k]$, where: $\bar{\lambda}[d] \in \Psi_p[d]$; $\lambda[d,k] \in \theta_{\mathrm{exog}}$; $\bar{\lambda} \in \Xi^D$; and $\lambda[d] \in \Xi^K$ for all $d$.

**R2e.** Outside an experiment $\theta_{\mathrm{clock}}$ and $\theta_{\mathrm{exp}}$ are single valued and have no affect on $U(\alpha, \theta)$ and $PP\{\theta' \,|\, \alpha, \theta\}$.

**R2f.** The exogenous vector has length $P = K(D+N)$ and takes on exactly one value,

defined in three ways:

$$\theta_{\text{exog}} \equiv (\ \lambda[1,1] \quad \cdots \lambda[D,K] \quad \Upsilon_1 \quad \Pi_1 \quad \delta_1 \quad \rho_1 \quad \cdots \quad \Upsilon_K \quad \Pi_K \quad \delta_K \quad \rho_K \ )$$

$$\equiv (\ \Lambda[1] \quad \Lambda[2] \cdots \quad \Lambda[D] \quad \Gamma[1] \quad \cdots \quad \Gamma[K] \ )$$

$$\equiv (\ \Lambda \quad \Gamma \ ). \tag{9}$$

The bounded open set of possible exogenous vectors is:

$$\hat{\Theta}_{\text{exog}} \equiv \{\Xi^K\}^D \times \left\{ \mathbf{U} \times \mathbf{P} \times (0,1)^2 \right\}^K \subset \Re^P. \tag{10}$$

The exogenous vector contains all parameters of the model to be estimated from data. The ˆ is included in $\hat{\Theta}_{\text{exog}}$ to avoid confusing the set of valid estimates of $\theta_{\text{exog}}$ and the set of values that the exogenous vector takes on in a social experiment, which is by definition the singleton $\{\theta_{\text{exog}}\}$. If the parameters $\Gamma$ and proportions $\Lambda$ are estimated freely, then complicated correlations between policy parameters and demographic variables are possible, and membership in group $d$ can be considered a lagged endogenous choice. A policy experiment that changes values in $\theta_{\text{pol}}$ without changing $\bar{\Lambda}$ or $\Lambda$ assumes that the policy change would not invoke a response in demographic variables. However, mixing probabilities can be adjusted along with policy parameters if there is other evidence on how $d$ would respond to changes in $\theta_{\text{pol}}$.

The primitive outcome-to-state transition and the endogenous choice probabilities combine to generate the state-to-state transition $P_s$ defined in the Appendix. By requiring that $k$ and $d$ are indices for all invariant (non-ergodic) endogenous variables, a stationary (or ergodic) distribution over the state space is determined by $P_s$, $\bar{\Lambda}$, and $\Lambda$. This stationary distribution, denoted $P_{-\infty}(\theta)$, describes the population prior to the experiment.

**Implication 1: Ergodic and smooth probabilities.** Under R1-R2:

**I1a.** There exists a unique distribution $P_{-\infty}\{\theta\}$ such that

$$\forall\, d, k, \qquad \bar{\lambda}[d]\lambda[d,k] = \sum_{\theta'} \mathcal{B}\left[k = k', d = d'\right] P_{-\infty}\{\theta'\} \tag{11}$$

$$\forall\, \theta \in \Theta, \qquad P_{-\infty}\{\theta\} = \sum_{\theta'} P_s\left\{\theta \mid \theta'\right\} P_{-\infty}\{\theta'\}. \tag{12}$$

**I1b.** $P\{\alpha\,|\,\theta\}$, $P_s\{\theta'\,|\,\theta\}$ and $P_{-\infty}$ are smooth functions of $\theta_{\mathrm{exog}}$ on $\hat{\Theta}_{\mathrm{exog}}$.

All proofs are provided in the Appendix.

These results are nearly the same as Rust's (1994) smoothness results, except in three respects. Rust specifies an error term (often extreme value) in the utility with infinite support to ensure smooth choice probabilities. Instead, the approach followed by Eckstein and Wolpin (1999) and others is followed here by smoothing choice probabilities without an error term in the utility. The state space for the subject's problem therefore remains discrete (and bounded). Second, Rust assumed a 'transparent' environment (defined below) and was not concerned with sampling from an endogenous distribution over the state space. Third, in Rust the observed component of the state vector could follow an arbitrary transition that could include absorbing states. Here, extra structure is placed on the transition to guarantee an ergodic distribution.

## II.C Measurements

From a subject's point of view the outcome $(\alpha, \theta)$ completely describes their current situation and what can be known about the future. Realistically the experimenter cannot observe all the aspects of an individual's current outcome. Furthermore, social experiments are carried out on selected populations, and selection is based on observed characteristics of the subjects. A description of what is observed (or measured) by the experimenter is made an integral part of the model. Measurements are also called *results*.

**Requirement 3. Measurements.**

**R3a.** $Y : A(\Theta) \times \Theta \to \Re^M$ is the measured result vector generated by outcomes. $Y(\alpha, \theta)$ does not include $k$ or $\theta_{\mathrm{exog}}$.

**R3b.** The state variables in the vectors $\theta_{\mathrm{obs}}$ and $\theta_{\mathrm{cond}}$, defined as

$$\theta_{\mathrm{obs}} \equiv (\,\theta_{\mathrm{exp}} \quad d\,) = (\,g \quad e \quad d\,) \tag{13}$$

$$\theta_{\mathrm{cond}} \equiv (\,t \quad \theta_{\mathrm{obs}}\,) = (\,t \quad g \quad e \quad d\,), \tag{14}$$

10

are by definition observed but excluded from $Y(\alpha, \theta)$.

**R3c.** Sample information is contained in $y = [y(\theta_{\text{obs}})]$, where $y(\theta_{\text{obs}})$ is a vector, $y(\theta_{\text{obs}}) = [y_n(\theta_{\text{obs}})]$, and where $y_n(\theta_{\text{obs}})$ is a panel of measurements for subject $n$ in measurement group $\theta_{\text{obs}}$. That is, $y(\theta_{\text{obs}}) = [y_n(\theta_{\text{cond}})]$ and $y_n(\theta_{\text{cond}})$ corresponds to $Y(\alpha, \theta)$ for subject $n$ in period $t$.

**R3d.** The number of subjects in each group is denoted $N(\theta_{\text{obs}})$. Missing information is denoted $\dot{y}$. The number of observations available in a conditioning period is denoted $N(\theta_{\text{cond}}) = \sum_{n=1}^{N(\theta_{\text{obs}})} \mathcal{B}[y_n(\theta_{\text{cond}}) \neq \dot{y}]$.

The expected measured result conditional on arriving at state $\theta$ is

$$E[Y \mid \theta] \equiv \sum_{\alpha \in A(\theta)} P\{\alpha \mid \theta\} Y(\alpha, \theta). \tag{15}$$

Since there is an ergodic distribution, there is a constant expected outcome within each observed group. That is,

$$E_{-\infty}[Y \mid d] \equiv \sum_{\theta'_{\text{end}}} \mathcal{B}[d = d'] E[Y \mid \theta'] P_{-\infty}\{\theta'\}/\bar{\lambda}[d] \tag{16}$$

is the mean status quo result before an experiment is carried out. It is not necessary to sum over elements of the state vector outside of $\theta_{\text{exog}}$ because the requirements so far have ensured that in the social environment they are either single-valued or have no effect on subject behavior.

**Definition 2. Social Environment.** A social environment

$$\mathcal{V}(\,\Theta,\, \mathbf{A}(\theta),\, U(\alpha,\theta),\, P\{\theta' \mid \alpha,\, \theta\},\, Y(\alpha,\theta),\, y\,)$$

includes a state space, action sets, a utility, a transition, a measurement, and data that satisfy R1-R3.

**Examples.** A social environment $\mathcal{V}$ is

◇ *autonomous* if no subject can affect the future: $\forall(\alpha, \theta),\, P\{\theta' \mid \alpha,\, \theta\} = P\{\theta' \mid \theta\}$;

◇ *myopic* if no subject places value on future states: $\forall\, k,\ \delta_k \approx 0$;

◇ *static* if it is autonomous or myopic;

◇ *irrational* if $\forall\, k,\ \rho \approx 0$;

◇ *homogeneous* if $K = 1$;

◇ *transparent* if $\alpha \in Y(\alpha, \theta)$ and $\theta_{\text{end}}\big|_{\neg k, d} \in Y(\alpha, \theta)$.

In a myopic environment subjects do not care about the future, while in an autonomous environment they may care about the future but cannot influence it. In either case the future is irrelevant to decisions made today so $P_a\{\,\alpha\,|\,\theta\,\}$ is a function of $U(\alpha, \theta)$ alone and we might call the environment static. In an irrational environment choice probabilities are equal for all feasible choices, so measured results are unrelated to the value of actions.

In a homogeneous environment, subjects in the same demographic group have identical problems and identical conditional choice probabilities. In a transparent environment the observer has the same information about the distribution over next period's state as the subject, except for the permanent values $k$ and $\theta_{\text{exog}}$. A weaker assumption (not listed above) is conditional homogeneity. In this case subjects that share some observed qualities, including perhaps measured outcomes before random assignment, also share unobserved type with high probability. This assumption relates to matching estimators described in Heckman et al. (1999) and the assumption of 'unconfoundedness' in Hotz et a. (1999). That is, in a conditionally homogeneous environment, pre-assignment outcomes can be used to control for endogenous selection of subjects into a non-experimental program.

# III. The Experiment

## III.A Treatment

An experiment introduces into a social environment a program of treatment that takes place in a finite number of phases each with a finite maximum duration. A subject's status in the program is described by the sub-vector $\theta_{\text{clock}}$ in (3). There are three required clock variables, the phase of treatment, $f$, the time spent residing in the phase so far, $r$, and the experimental period, $t$, at which measurements for the current outcome will be made. Before a subject enters treatment they are in phase $f = 0$, which corresponds to the 'real world' of the social environment. After completing treatment a subject enters the last phase $f = F$, which in some cases is the same as phase 0. Within a phase, $r$ determines the future course of treatment but not current utility. A phase between 0 and $F$ lasts at most $R[f]$ periods. The experimental time $t$ does not affect a subject and is used to coordinate outcomes across experimental groups.

The results for a subject determine the course of treatment, and based on the subject's information the clock setting next period is deterministic. Since all new state transitions created by the experiment are excluded from the social environment $\mathcal{V}$, the model requires that all effects of treatment enter through adjustments to the utility or transition as functions of the experimental clock. Parameters of the treatment that may be subject to hypothetical variation are collected in a vector $\Psi_t$.

**Requirement 4. Program of Treatment.**

**R4a.** The treatment vector takes the form:

$$\Psi_t = \begin{pmatrix} R & f_+(y; \theta_{\text{clock}}) & \cdots \end{pmatrix} \tag{17}$$

where $R$ is a $F - 1$ vector of positive integers denoting the maximum length of phases $f = 1, 2, \ldots, F - 1$, and $f_+ (y; \theta_{\text{clock}})$ is the phase increment next period. $U(\alpha, \theta)$ and $P\{\theta' \,|\, \alpha, \theta\}$ are functions of the remaining (modeler-specified) elements of $\Psi_t$, which can interact with the current phase $f$ and treatment group $g$.

13

**R4b.** Treatment is deterministic: $\forall\,(\alpha,\theta):0<f<F$,

$$f' = f^\star\left(f+\mathcal{B}\left[r=R[f]\right]+f_+,0,\emptyset\right) \tag{18}$$

$$r' = r^\star\left(\mathcal{B}\left[f_+=0,r<R\left[f\right]\right]r+1,0,\emptyset\right) \tag{19}$$

$$t' = t^\star(t+1,0,\emptyset). \tag{20}$$

**R4c.** Treatment is progressive: $f\le f+f_+\le F$ and $f<f+f_+$ when $r=R\left[f\right]$.

**R4d.** Utility and the transition of state variables (other than those in $\theta_{\mathrm{clock}}$) are unaffected by $r$ and $t$: $\forall\,(\alpha,\theta),q\in\theta_{\mathrm{endog}},\tilde{r},\tilde{t}$,

$$U(\alpha,\theta)=U(\alpha,\theta)\Big|_{\tilde{r},\tilde{t}} \tag{21}$$

$$q^\star(\bar{q},[\pi_j],[\mathbf{Q}_j])=q^\star(\bar{q},[\pi_j],[\mathbf{Q}_j])\Big|_{\tilde{r},\tilde{t}}. \tag{22}$$

The transition rules for treatment are simpler than the notation might suggest. The phase next period can only depend on $\theta_{\mathrm{clock}}$ and the measurement vector $Y(\alpha,\theta)$, because the experimenter must be able to set the phase based on observed results. If a subject has not reached $R[f]$, their next phase can be the current phase ($f_+=0$) or some other further phase ($f_+>0$), including the post-treatment phase ($f_+=F-f$). When $r=R[F]$ the only difference is that the default phase becomes $f+1$ rather than $f$. The discrete jump process for $r$ simply says that $r'=r+1$ if the phase next period will be the same as this period. Otherwise $r=1$. Since the treatment program is deterministic given measurements and group assignment, it follows that $r$ and $f$ can be determined for a subject based on the data generated, $y_n\left(\theta_{\mathrm{cond}}\right)$.

**Examples of treatment.**

A phase $1\le f<F-1$ is

⋄ a *qualification phase* if $\forall\,(\alpha,\theta)$, $U(\alpha,\theta)=U(\alpha,\theta)\Big|_F$ and $\forall\,q\in\theta_{\mathrm{end}}$, $q^\star()=q^\star()\Big|_F$;

⋄ *autonomous* if $f_+=0$;

⋄ a *waiting period* if it is an autonomous qualification phase;

14

- ◇ an *entry period* if it is a waiting period and $f = 1$;

- ◇ *voluntary* if $\forall\, \theta\; \exists\, \alpha \in A(\theta):\; f + \mathcal{B}\left[r = R[f]\right] + f_+ = F$.

In a qualification phase, utility is untreated, although the course of treatment may depend on results. In an autonomous phase the course of treatment does not depend on results. Combining these two properties results in a waiting period, and a waiting period at the beginning of an experiment is an entry period. The program of treatment does not specify how subjects enter treatment. Nor does it describe what happens once treatment ends by entering phase $F$. These are elements of an experiment associated with treatment groups which are defined next.

### III.B Experimental Groups and Policy Innovations

Each subject is associated with an experimental status described by the sub-vector $\theta_{\mathrm{exp}}$ introduced in (3). A single experiment might include multiple treatment groups and multiple criteria for entry. For example, subjects may be assigned different doses of the treatment or face different selection criteria. Differences across groups must be accounted for by either treatment status ($g$) or initial conditions ($e$). The control group is by definition $g = G$. In a simple experiment $G = 1$ and the treatment group is designated $g = 0$. When $G > 1$ more than one treatment group share the same control group. Upon random assignment the control group transits directly to $f = F$. Treatment groups begin with an initial clock setting $\bar{\theta}_{\mathrm{clock}}(\theta_{\mathrm{exp}})$.

Experimental (or measurement) time $t$ has to be synchronized across experimental groups because random assignment may occur after different selection criteria. Sample selection in a group $e$ occurs in a range of periods, $[t_{\min}(e), t_0(e)]$. The length of the selection period $T(e) = t_0(e) - t_{\min}$ is specified by the modeler. Random assignment occurs at the end of $t_0$, which is normalized to $0$ in one group and is less than zero for all other groups. Post-assignment measurements are made in the range $[t_0 + 1, t_{\max}]$. Because selection

15

occurs before assignment to treatment, the data process begins before treatment starts, before subjects are aware that they will take part in an experiment.

The modeler defines which results are feasible for each conditioning point in the experiment, $\theta_{\mathrm{cond}}$ defined in (14). Valid measurements are indicated with a Boolean function $\mathcal{H}[y; \theta_{\mathrm{cond}}]$. A non-valid measurement implies selection out of the sample, and prior to random assignment lead to ineligibility for the experiment. The criteria may differ across experimental groups, but not between treatments and controls for the same treatment. For example, suppose treatment group $e$ requires six periods of non-employment to be eligible for random assignment. Then $T(e) = 6$. If $z$ is an indicator for employment then $t \leq t_0$ implies $\mathcal{H}[y; \theta_{\mathrm{cond}}] = 1 - \mathcal{B}[z]$. Note that the state space must be constructed so that the selection criteria can be represented as a sequence of feasible results that subjects must satisfy. In this example $z$ must be an element of $y$ or directly inferable from it. The state space need not be expanded so that the complete criteria be computable from $y$. In this case, for example, it is unnecessary to include five lagged values of $z$ so that eligibility for the experiment can be deduced from a single result vector $y$. Post-assignment histories may also be selected if subjects can be lost due to attrition.

Subjects who meet eligibility requirements for group $e$ are assigned a value of $g$ according to a discrete jump process, $g = g^\star(g, [\mu_g], [g])$. Recall that the arguments of $g^\star$ can depend upon the current outcome $(\alpha, \theta)$. Thus, we can require that for $t \neq t_0$ group assignment is invariant: $\sum \mu_g = 0$. Jump probabilities are specified only apply at $t = t_0$. In a true experiment the probability $\mu_g$ is under the control of the experimenter, which implies that $\mu_g$ can only be a function of $d$, $e$, and the result vector $Y(\alpha, \theta)$. A 'randomized experiment' $\mu_g$ only depends on $e$ and would thus equal the sample proportions of each treatment gorup. At the other extreme, group assignment is under the direct control of the subject. The value of $g$ would be directly determined by the action vector $\alpha$. Note, however, that the subject does not get to choose the date of assignment $t_0$.

True experiments are usually of interest because they inform policy decisions. It is

16

presumed that the question,"How would the treatment affect the population if it were policy?" will be asked of any experiment. This framework distinguishes between an experiment and an unexpected change in policy, which will be called a policy innovation. An experiment is a temporary one-time application of a treatment to a selected population. Potential subjects unwittingly become eligible for treatment by making choices under the status quo. Treatment will end and all subjects return to the status quo environment. A policy innovation is a permanent application of a treatment to all eligible members of the population. The population is informed of the program of treatment and can enter treatment at any point they become eligible, either now or in the future.

The case $e = 0$ is reserved for policy innovations. A policy innovation is a new social environment that embeds the treatment into the status quo. A policy innovation has a short-run, non-stationary effect while the treated choice probabilities operate on the untreated state distribution generated by the old regime. When considering long-run effects, there is a conflict between the definitions of a social environment and a program of treatment. Treatment programs are of a finite length. Social environments are infinite. A finite treatment has no long-run impact in an ergodic environment. In contrast, when a treatment program is turned into a policy members of the population can expect to enter treatment any time they satisfy some criteria. To embed a finite treatment in an infinite social environment, policy innovations will be *recurring* treatments. A member of the population that completes treatment enters phase $F$ and then becomes eligible for treatment again by moving back to phase $0$. For simplicity, the transition from $F$ to $0$ occurs with a constant probability $\tau(g)$. The subject moves to phase 1 if they satisfy the criteria for treatment. This restores stationary to the environment while retaining the finite nature of treatment.

## Requirement 5. Experimental Groups

**R5a.** The parameters for experimental groups are organized as sub-vectors of $\Psi_x$,

$$\Psi_x(\theta_{\exp}) = \begin{pmatrix} T & \mathcal{H}\left[y;\theta_{\text{obs}}\right] & g_0^\star() & \bar{\theta}_{\text{clock}} & \tau \end{pmatrix}, \tag{23}$$

where: $T(e)$ is the length of the selection period before random assignment; $\mathcal{H}\left[y;\theta_{\text{cond}}\right] \in \{0,1\}$ is an indicator for whether $y$ is a feasible measurement at point $\theta_{\text{cond}}$; $g_0^\star()$ is the discrete jump process for $g$ at random assignment; $\bar{\theta}_{\text{clock}}\left(\theta_{\exp}\right)$ is the clock setting just after random assignment; and $\tau(g)$ is the treatment recurrence probability.

**R5b.** For $g = G$, $\bar{\theta}_{\text{clock}} = \begin{pmatrix} t_0 + 1 & 1 & F \end{pmatrix}$. Based on $T(e)$ and other values of $\bar{\theta}_{\text{clock}}$, observations are split into pre- and post-assignment periods, $t_{\min}(e) \le t_0(e) \le 0 \le t_{\max}(e)$, as described in the Appendix. Entry group assignment takes an iid discrete jump at $\bar{t} \equiv \min_e t_{\min}(e) - 1$; group assignment takes place at $t_0$:

$$E > 0, t = \bar{t} \;\rightarrow\; e' = e^\star(e, 1/E, [e]) \tag{24}$$

$$t \ne t_{\min} \;\rightarrow\; e' = e^\star(e, 0, \emptyset) \tag{25}$$

$$t \ne t_0 \;\rightarrow\; g' = g^\star(g, 0, \emptyset) \tag{26}$$

$$t = t_0 \;\rightarrow\; g' = g_0^\star(). \tag{27}$$

**R5c.** Given $e$, no selection occurs before $t_0$ and feasible measurements before random assignment are the same in all treatment groups:

$$t < t_{\min}(e) \;\rightarrow\; \mathcal{H}\left[y;\theta_{\text{cond}}\right] = 1$$
$$t <= t_0(e) \;\rightarrow\; \mathcal{H}\left[y;\theta_{\text{cond}}\right] = \mathcal{H}\left[y;\theta_{\text{cond}}\right]\Big|_{g'} \tag{28}$$

**R5d.** Pre-treatment status recurs with probability $\mathcal{B}\left[e = 0, g < G\right]\tau(g)$:

$$f = F \;\rightarrow\; f' = f^\star\left(F, \mathcal{B}\left[e = 0, g < G\right]\tau(g), 0\right). \tag{29}$$

For $e > 0$, a transition out of phase $0$ is never expected: $f = 0 < e \;\rightarrow\; f' = f^\star(f, 0, \emptyset)$. For $e = 0$, the transition occurs when the subject meets the earliest eligibility standard of any other entry group (see the Appendix).

18

A social experiment is simply a collection of one or more treatment and entry groups applied to a social environment.

**Definition 3. Social experiment.** A social experiment $\mathcal{E}(\Psi_x, \Psi_t, \mathcal{V})$ is a program of treatment and a set of experimental groups applied to a social environment $\mathcal{V}$ that satisfy R1-R5.

**Examples.** A treatment group or experiment is

⋄ *controlled* if choice probabilities in phase $0$ and phase $F$ are equivalent: $\forall (\alpha, \theta)$,

$$P\{\alpha | \theta\}\Big|_{f=0} = P\{\alpha | \theta\}\Big|_{F};$$

⋄ *selected* if $T > 0$ and there exists $y$ and $\theta_{\mathrm{cond}}$ such $\mathcal{H}[y; \theta_{\mathrm{cond}}] \neq 1$;

⋄ *randomized* if at $t = t_0$, $g^\star$ is an iid discrete jump process;

⋄ *a policy innovation* if $e = 0$, and is *recurring* if $\tau > 0$;

⋄ *natural* if it is a non-recurring policy innovation and $g$ corresponds to subsets of $d$: $\forall g \, \exists \, \mathbf{D}_g \subseteq \{0, 1, \ldots, D\} : t = t_0, d \in \mathbf{D}_{\tilde{g}} \; \rightarrow \; g' = g^\star(\tilde{g}, 0, \emptyset)$;

⋄ *naturally randomized* if it is natural and for all $d \in \mathbf{D}_g$, there exists a $d' \in \mathbf{D}_G$: $\Psi_p[d] = \Psi_p[d']$ and $\Lambda[d] = \Lambda[d']$.

An experiment is a surprise change in the environment which alters rational behavior from that point on. A policy innovation is a surprise change in the environment that becomes part of the environment. The difference between a true experiment ($e > 0$) and a policy innovation ($e = 0$) is that a subject of an experiment has only one unanticipated chance at treatment. If subjects of a true experiment are not eligible for treatment when the experiment occurs then they are never eligible again. In a policy innovation an individual knows that the treatment will be available (or mandatory) in the future. The difference between a recurring ($\tau > 0$) and non-recurring ($\tau = 0$) innovation is that in a recurring innovation receipt of treatment does not preclude receiving it later on.

In a controlled experiment those assigned to the control group act exactly as they would have if they had not been selected from the population in the first place. They also act (conditional on the current state) as the treatment groups will act after treatment ends. This rules out Hawthorne and placebo effects. R5.3 and randomization ensure that the distribution over states at random assignment are identical across treatment groups, but not necessarily across experimental groups due to differences in selection criteria. An experiment may be controlled but not randomized if $g$ is correlated with endogenous states. Since $\theta_{\text{exog}}$ contains $d$ and $k$, randomization requires that assignment to $g$ is uncorrelated with both observed and unobserved types. Independence is required just prior to applying the first selection criterion for $e$.

In all types of experiments phase 0 and phase F have infinite horizons. With $\tau = 0$ an individual knows that they leave phase 0 and enter treatment only once. Since treatment is progressive they eventually enter the absorbing phase $F$, which in this case is the environment before the innovation or experiment. With $\tau > 0$ the previous environment is replaced altogether because in phase $F$ the positive probability of undergoing treatment again affects the value of choices. The value of choices in phase $F$ affect the value of choices in $f = 0$, and vice versa. The value function and choice probabilities in these two phases must be computed simultaneously in all phases. If $S$ denotes the size of the endogenous state space (the number of distinct values of $\theta_{\text{end}}\big|_{\neg d,k}$), then $S$ is the size of the state space for a single subject in the status quo environment. A recurring policy innovation expands the state space for a single subject to $S \times S \times \prod_{f=1}^{F-1} R[f]$.

**Implication 2: Natural experiments as controlled and randomized experiments.**

  **I2a.** A natural experiment is not randomized.

  **I2b.** A natural experiment conducted on a myopic environment is controlled.

Prior to random assignment subjects of a policy innovation can anticipate assignment and would (presumably) make different decisions than in the post-treatment phase. In a

dynamic environment there is no way to guarantee that a natural experiment is controlled without specifying more aspects of the social environment. In a myopic environment subjects do not anticipate the future so even if they are informed of the policy innovation they will act the same before random assignment as after any treatment ends. Technically, the definition of randomization cannot be satisfied by a natural experiment. Thus a naturally randomized experiment is defined as a natural experiment where assignment to treatment group is based on an demographic variable that is irrelevant to the policy vector.

There are a two types of situations often referred to as natural experiments which, in this framework, would be classified as a true experiment ($E > 0$) or not as an experiment at all. First, the *temporary* expiration of a policy that one can credibly model as unanticipated in the status quo environment. Although this situation is not designed it meets all the other requirements of a deliberate experiment. Such a situation is controlled if there exists some demographic group not affected by the temporary expiration but which is modelled as identical (in a policy sense) as the affected group. On the other hand, an ordinary policy innovation such as raising the minimum wage does not fit as an experiment in this framework. As with a natural experiment, assignment to treatment is based on a demographic characteristic. Subjects enter 'treatment' by being affected by the new minimum wage. Unless the change is temporary agents in the treatment group will modify behavior permanently even in the status quo. In addition, outcomes in the post-treatment phase are dependent on treatment since subjects can expect to be affected by the minimum wage at some future date. Within this framework, a permanent change in policy should be modelled as a recurrent policy innovation: $E = 0$ and $\tau > 0$. Under conditions on the social environment similar to those for a natural experiment, such a policy innovation can still end up being controlled.

### III.C Examples of Social Experiments

21

The Appendix shows how four situations can be placed within this framework: the Illinois Re-employment Bonus Experiment, the NIT Experiments, the Vietnam Draft Lottery natural experiment, and non-experimental lifecycle data.

## IV. Predicting Experimental Results

### IV.A Sample Selection

The model of subject behavior generates a stationary distribution across states outside the experiment, denoted $P_{-\infty}\{\theta\}$ and defined in (12). The data generating process begins with the the experimenter setting $t = \bar{t}$ for all points in the state space and then letting the ordinary transitions take place. The distribution across states at the start of $\bar{t} + 1$, given type $k$, is

$$\Omega\{\,\theta' \mid k, \theta_{\text{cond}}\} \equiv \frac{1}{\max\{E, 1\}\bar{\lambda}[d]\lambda[k, d]} \sum_{\theta} P_s\{\theta' \mid \theta\}\,\mathcal{B}[k]\,P_{-\infty}\{\theta\}\,\Big|_{\bar{t}}. \tag{30}$$

Requirement R5.2 and the presence of $E$ (the number of experimental groups) in the denominator ensure that each experimental groups is a copy of the target population. Define $\omega\,(k; \theta_{\text{cond}})$ as the cumulative proportion of type $k$ that has survived selection up to the start of period $t$ in conditioning group $\theta_{\text{cond}}$. For all $k$ this proportion is initialized as

$$\omega\,(k; \theta_{\text{cond}})\,\Big|_{\bar{t}} \equiv 1. \tag{31}$$

For computational purposes, $\omega\,(k; \theta_{\text{cond}})$ is *not* defined as the distribution of $k$ given $\theta_{\text{cond}}$. This allows predictions to be computed in parallel for all values of $k$. Only when predictions are confronted with data will the distribution across $k$ be 'fixed up' using $\omega\,(k; \theta_{\text{cond}})$ and $\lambda[d, k]$.

At $t = \bar{t} + 1$ the selection process for at least one group begins. In these groups the distribution in the next period will condition on satisfying the first entry condition defined by $\mathcal{H}\,[y; \theta_{\text{cond}}]$. When $t$ reaches $t_0$ in a group, random assignment to treatment group $g$ and

22

the initial clock setting takes place between observation of the current result and the start of period $t_0 + 1$. The subject 'wakes up' in period $t_0 + 1$ either undergoing treatment or in the control group. By assumption (and for simplicity) they immediately re-optimize given their initial state and the program of treatment, including the case that a recurring policy innovation ($E = 0$, $\tau > 0$) has occurred and their environment has changed forever. The generalized transition that accounts for selection (feasible histories) and random assignment is written

$$P^\star\left\{\theta' \mid \theta\right\} = \sum_\alpha P\left\{\theta'\Big|_{\mathcal{B}[t=t_0]\bar{\theta}_{\text{clock}}} \mid \alpha, \theta\right\} \mathcal{H}\left[Y(\alpha,\theta); \theta_{\text{cond}}\right] P\left\{\alpha \mid \theta\right\}. \tag{32}$$

The assignment $\mathcal{B}\left[t = t_0\right] \bar{\theta}_{\text{clock}}\left(\theta_{\text{exp}}\right)$ is shorthand for saying that the insertion of the initial clock setting for $\theta'$ occurs only when making the transition from $t_0$ to $t_0 + 1$. (The clock setting depends on the new value of $g$ which cannot be determined from old outcome when $g$ is not deterministic. Therefore, initializing the clock cannot be described as a discrete jump process.)

In general, the conditional distribution of subjects across $\theta$ at the start of any period $t \geq t_{\text{min}}$ is denoted $\Omega\left\{\theta \mid k, \theta_{\text{cond}}\right\}$. Subjects will choose actions that determine their outcome $(\alpha, \theta)$ that period. The experimenter measures the outcome $Y(\alpha, \theta)$ if the result is feasible. If so, the subject's state next period is realized and they will contribute to the distribution next period. Otherwise they leave the sample and the distribution at $t+1$ must be corrected for the proportion of outcomes for which $\mathcal{H}\left[y; \theta_{\text{cond}}\right] = 0$. The cumulative proportion of the type $k$ population making it to $t$ is

$$\omega\left(k; \theta_{\text{cond}}\right) = \omega\left(k; \theta_{\text{cond}}\Big|_{t-1}\right)\left[\sum_{\theta'}\sum_{\theta} P^\star\left\{\theta' \mid \theta\right\} \Omega\left\{\theta \mid k, \theta_{\text{cond}}\Big|_{t-1}\right\}\right]. \tag{33}$$

The sequence $\omega\left(k; \theta_{\text{cond}}\right)$ is non-increasing in $t$. The distribution across states at any $t$ can now be defined recursively in terms of the $t-1$ distribution:

$$\Omega\left\{\theta' \mid k, \theta_{\text{cond}}\right\} = \frac{\omega\left(k; \theta_{\text{cond}}\Big|_{t-1}\right)}{\omega\left(k; \theta_{\text{cond}}\right)} \sum_{\theta} P^\star\left\{\theta' \mid \theta\right\} \Omega\left\{\theta \mid k, \theta_{\text{cond}}\Big|_{t-1}\right\}. \tag{34}$$

Note that for $t \leq t_0$ this distribution only describes subject's still (partially) eligible for the sample. The selection in periods until the end of $t_0$ have not yet been imposed. The distribution of the invariant types in an experimental group (at time $t$) is:

$$\lambda^\star(k|\theta_{\mathrm{cond}}) \equiv \lambda[k,d] \frac{\omega\left(k; \theta_{\mathrm{cond}}\right)}{\sum_{k'=1}^{K} \omega\left(k'; \theta_{\mathrm{cond}}\right)}. \tag{35}$$

At $t_{\mathrm{min}} - 1$ this would equal $\lambda[k,d]$, *if* the experiment is randomized. Once selection begins, the joint distribution of $k$ and $d$ drifts away from the population proportion, because subjects of differing types have different propensities to satisfy the selection criteria in $\theta_{\mathrm{exp}}$. If selection ends at $t_0$, then $\lambda^\star$ becomes constant for all $t > t_0$:

$$t > t_0 \;\; \rightarrow \;\; \lambda^\star\left(k|\theta_{\mathrm{cond}}\right) = \lambda^\star\left(k|\theta_{\mathrm{cond}}\big|_{t_0+1}\right) \equiv \lambda_0^\star\left(k|\theta_{\mathrm{obs}}\right). \tag{36}$$

With attrition, the condition on $t$ would change to the period at which attrition ends.

**Implication 3: Solution algorithm.** The model of subject behavior and empirical results in a social experiment can be solved by the algorithm in the Appendix.

### IV.B Mean Outcomes and Mean Treatment Effects

In an experiment with demographic variation, multiple outcomes and treatments, and repeated observations, impact is a vector conditioned on a subset of the state vectors $\theta_{\mathrm{cond}}$ introduced in (14). These variables are, by definition, exogenous to individual outcomes, conditional on pre-assignment behavior satisfying the sample-selection criteria. This section derives the conditional mean outcomes in a social experiment and derives their relation to individual outcomes and individual impact.

The mean observed outcome and the empirical impact in treatment group $g$ are

$$\hat{E}\left[Y \mid \theta_{\mathrm{cond}}\right] = \frac{1}{N(\theta_{\mathrm{cond}})} \sum_{n=1}^{N(\theta_{\mathrm{cond}})} \mathcal{B}\left[y_n(\theta_{\mathrm{cond}}) \neq \dot{y}\right] y_n(\theta_{\mathrm{cond}}) \tag{37}$$

$$\hat{\Delta}(\theta_{\mathrm{cond}}) \equiv \hat{E}\left[Y \mid \theta_{\mathrm{cond}}\right] - \hat{E}\left[Y \mid \theta_{\mathrm{cond}}\big|_G\right], \tag{38}$$

24

where $N(\theta_{\text{cond}})$ was defined in R3. Using (34), the model analogue to the empirical values are

$$E[Y \mid \theta_{\text{cond}}] \; = \; \sum_{\theta} \lambda^{\star}\left(k|\theta_{\text{cond}}\right) \Omega \left\{\, \theta \mid k, \theta_{\text{cond}} \,\right\} E\left[Y|\theta\right] \tag{39}$$

$$\Delta\left(\theta_{\text{cond}}\right) \; \equiv \; E\left[Y \mid \theta_{\text{cond}}\right] - E\left[Y \mid \theta_{\text{cond}}\Big|_{G}\right]. \tag{40}$$

The impact in (40) varies over experimental time $t$ and experimental group $e$ for several reasons. First, the treatment effect itself is non-stationary over $t$ as the subjects progress through treatment. Given $g$ the treatment is constant, but subjects in different experimental groups entered treatment at different stages. Selection into the experiment also creates a non-stationarity in both treatment and control groups as the distribution across states drifts back to the unselected distribution $P_{-\infty}\{\theta\}$. That is, if the experimental group $e$ is not a random sample of the population then the distribution over states in the control group $G$ varies with experimental time $t$ even though the controls never leave the status quo environment.

It would seem desirable that a particular treatment should have an impact that is not dependent on the experimental design. Otherwise, there would be little hope of using the experimental results to make policy-relevant statements outside of the experiment. The status quo outcome $E_{-\infty}\left[Y \mid d\right]$ defined in (16) can stand for the long-run analogue to $E[Y_1]$. It might appear straightforward to define the analogue to $E[Y_0]$ as well, but it is not. The first attempt might be to let the non-stationarity in $t$ die out by looking at outcomes as $t \to \infty$. After selection ends the distribution over states begins to converge to an ergodic distribution. This distribution is not the same as $P_{-\infty}\{\theta\}$ because the unobserved type $k$ is also selected, and its distribution is not ergodic. The long-run post-assignment distribution in group $\theta_{\text{exp}}$ is a re-weighted version of the ergodic distribution for the selection-bias on unobserved types.

**Implication 4: Long Run Impact.**

**I4a.** For $e > 0$ or $\tau = 0$,

$$\forall \theta, \quad P_{+\infty}\{\theta \mid \theta_{\text{cond}}\} \equiv \lim_{t \to \infty} \lambda_0^\star(k|\theta_{\text{obs}}) \, \Omega\{\theta \mid k, \theta_{\text{cond}}\} = \left( \frac{\lambda_0^\star(k|\theta_{\text{obs}})}{\lambda[k,d]} \right) P_{-\infty}\{\theta \mid d\}. \quad (41)$$

**I4b.** For $e > 0$ the long-run *in-sample* impact of an experiment or non-recurrent policy is independent of the treatment program, $\Psi_t$.

**I4c.** In a randomized experiment or in an experiment carried out on a homogeneous environment, the long-run in-sample impact is zero.

The long-run in-sample impact is not necessarily zero, but only because of selection on permanent unobserved heterogeneity. The true impact of a treatment on an individual can, of course, be *long-lasting* but not *permanent*. How long treatment lasts depends on many factors including the transition $P_S\{\theta' \mid \theta\}$, which determines how quickly a short-term perturbance gets washed away by the ergodic nature of the social environment.

The point is that in a selected experiment there is no hope of 'waiting out' the selection effects. Even if the effect of sample-selection on transient heterogeneity dies away quickly, the treatment effect dies away before selection on unobserved type does. At no point in an experiment does the selection component of empirical impacts give way to a pure treatment effect. On the other hand, suppose that there is enough confidence that the experiment was randomized that the assumption can be maintained. Then if, after a long period past the end of treatment, the empirical impact does not go to zero one could reject the assumption of a homogeneous environment.

To restate the problem: in defining a timeless impact of an experimental treatment applied to a social environment there is no way to disentangle non-stationary selection effects from non-stationary treatment effects when studying outcomes of an experiment. However, any social experiment also implies a version of the treatment $g$ in which it becomes policy. Requirement R4 codified this by reserving experimental group $e = 0$ for policy innovations.

**Definition 4. Policy Outcomes and Impact**

**D4a.** The long-run average effect of implementing treatment $g$ as a policy in demographic group $d$ is defined as

$$E_{+\infty}\left[Y \mid d, g\right] \equiv E\left[Y \mid \theta_{\text{cond}}\Big|_{e=f=0,g}\right]. \tag{42}$$

**D4b.** The long-run average policy impact of implementing treatment $g$ is

$$\Delta_{\infty}[d, g] \equiv E_{+\infty}\left[Y \mid d, g\right] - E_{-\infty}\left[Y \mid d\right]. \tag{43}$$

The notation and the substance behind (42) resolve both of the problems in using $P_{+\infty}$ to compute long-run impact. First, $E_{+\infty}\left[Y \mid d, g\right]$ does depend on $g$ because by definition a policy innovation is not a finite disturbance to the social environment, unless $\tau(g) = 0$. Second, $E_{+\infty}$ does not depend on the selection criteria, because all experimental groups are mapped into $e = 0$ for the purpose of calculating policy innovations. Selection bias is explicitly avoided by measuring outcomes in phase $f = 0$ before any selection criteria have been applied.

### IV.C Individual Treatment Effects

The fundamental problem that social experiments attempt to solve is that individual outcomes under treatment and control status are never both experienced by the same subject. Yet, as with mean impacts, defining individual outcomes is not straightforward in a social environment. As a first pass consider:

$$Y_g \overset{?}{\equiv} Y\left(\alpha, \theta\right)\Big|_g. \tag{44}$$

This reveals a poverty of notation when applying the notation of section II to the dynamic and uncertain social environment defined in section III. The theoretical result depends on the whole state, including possibly unobserved actions, state variables, policy parameters, and exogenous parameters. Clearly $Y_g$ must be considered a random variable defined on a sample space that includes realizations not available at all possible time periods. Perhaps

$Y_g$ is implicitly based on subjective information available at time $t_0$ or perhaps time $t_{\min} - 1$. A further complication arises when results are measured at a number of periods after random assignment. In this case the outcome $(\alpha, \theta)$ in (44) has date $t$, but at $t_0$ elements of the state at $t$ are still unrealized, even given the subject's information set at time $t_0$. Thus, the right-hand side of equation (44) is unrealized until $t$. Until then the outcome can only be an expectation based on an information set available in a certain period.

More precisely, $Y_g$ is supposed to be the answer to the question, "Given what you know now at time $\tilde{t}$, if you were placed in situation $g$ at time $t_0$, what do you expect the result to be at time $t$?" As always, $t$ is the date of outcome measurement, $t_0$ is the date of random assignment, and $\tilde{t}$ is the date when the expectation is taken. Presumably the information set does not include realizations after random assignment, because by then the subject is already reacting to being assigned to $g$. Hence, $\tilde{t} \leq t_0$. But it is also problematic to condition upon less information than before random assignment. This requires the subject to suspend knowing that they will be placed in group $g$ only at $t_0$. In other words, for $\tilde{t} \leq t < t_0$ the subject must answer as if they are pretending that they don't know random assignment will occur at $t_0$. Without this pretense the actual random assignment occurs at $\tilde{t}$ and the treatment is altered because the time between $\tilde{t}$ and $t_0$ becomes an auxiliary qualification phase. One notion of individual treatment avoids this problem by setting $\tilde{t} = t_0$.

**Definition 5. Individual Treatment Effects** The *expected subjective treatment effect* is the expected difference in results between assignment to a treatment group and a control group given the information available to the subject as of random assignment at the end of $t_0$:

$$\Delta_{\mathrm{S}}(\alpha^0, \theta^0) \equiv E\left[Y(\alpha, \theta) \mid (\alpha^0, \theta^0)\right] - E\left[Y(\alpha, \theta)\big|_G \mid (\alpha^0, \theta^0)\right] \tag{45}$$

Notice that in the atheoretic framework $\Delta = Y_0 - Y_1$. It would be tempting to consider associating $\Delta_S$ with $E[Y(\alpha, \theta) - Y(\alpha, \theta)\big|_g]$. Random assignment would appear to imply that

28

conditioning on assignment as of time $t_0$ has no effect on expectations, making it possible to interchange the expectation and subtraction operators in (45). However, this is valid only in special social experiments carried out on special social environments in which the distribution over states does not depend upon which group the subject is assigned to.

**Implication 5: Individual treatment as an expected difference.** When:

  **I5a.** $t = t_0 + 1$

  **I5b.** or $\Psi_t$ is autonomous **and** type $k$ is autonomous or irrational,

then

$$\Delta_{\mathrm{S}}(\alpha^0, \theta^0) = E\left[ Y(\alpha, \theta) - Y(\alpha, \theta)\Big|_G \mid (\alpha^0, \theta^0) \right].$$

Even if the social environment is autonomous the experimental treatment may not be. Since treatment may have a direct effect on measured results, the distribution over $\theta_{\mathrm{clock}}$ must also not be affected by realized outcomes between $t_0$ and $t$. In effect nothing subjects do or experience between time $t_0$ and time $t$ can alter the state distribution. In other environments and other experiments the subjective impact cannot be expressed as an expected difference in results. For example, a myopic environment ($\delta = 0$) is not sufficient for the result. Although myopic subjects place no weight on outcomes between now and $t$, their myopic choices between time $t_0$ and time $t$ alter the distribution across states at time $t$ when the environment and experiment are not autonomous. Thus, perhaps counter-intuitively, a static decision framework is neither necessary nor sufficient to identify individual treatment impact with $\Delta_{\mathrm{S}}$ in (45).

Although the difference in mean outcomes can be treated as an unbiased estimate of the mean over individual impacts in a timeless model, in other environments individual subjective impact as defined above fails to match up with empirical impact in a transparent way. Does there exist a notion of individual and subjective impact for which (38) is an unbiased estimate? The answer, obviously, is the notion in which individuals form expectations under the same information set as the experimenter uses to condition

29

empirical impact.

**Implication 6: Subjective treatment effects.** The mean treatment effect (38) corresponds to an expected individual effect (45) when:

    **I6a.** the environment is irrational or

    **I6b.** the environment is autonomous and homogeneous, or

    **I6c.** subjects conduct the following thought experiment: they set the probability of being in any state in the outside environment equal to $P_{-\infty}\{\theta \mid d\}$; they form the expectation conditional upon becoming eligible for the sample; but for $t_{\min} \leq t \leq t_0$ they set choice probabilities as if they were ignorant that random assignment is to occur at the end of $t_0$.

What happened to the straightforward and intuitively appealing individual impact $\Delta$ emphasized in the treatment effects literature? The problem is that these impacts condition on the value of one variable, $t$, that is irrelevant to the individual while at the same time integrating out information available to and pertinent to the individual (type, initial state, and experimental clock). Therefore differences between mean outcomes within groups cannot typically be interpreted as the average of treatment effects felt by individuals.

# V. Applications and Extensions

## V.A Out-of-Sample Prediction

An out-of-sample prediction is defined as using estimates within the same social environment $\mathcal{V}$ that surrounds the experiment. Outcomes can be predicted for hypothetical treatments and selection criteria by simply including them in $\theta_{\exp}$ with sample proportions equal to zero. Hypothetical experiments conducted in the same environment would be based on the same result function $Y(\alpha, \theta)$, and these predictions would have a well-defined confidence interval associated with them. As with any out-of-sample predictions,

the confidence interval widens with the distance between the hypothetical conditions and the in-sample conditions. In the case of a social experiment, the more different the hypothetical selection criteria and program of treatment are from those in the experiment, the less confidence would be placed on any predictions.

### V.B Out-of-Population Prediction

An out-of-population prediction uses the estimated model in a different social environment, denoted $\mathcal{V}^\perp$. For simplicity, consider a hypothetical social environment with a single demographic group ($D^\perp = 1$) that differs from the estimated environment in just four respects. The new demographic group faces its own policy vector $\theta_{\text{pol}}^\perp$; it has its own distribution of unobserved types $\Lambda^\perp$; it is measured using its own vector $Y^\perp$; and data are based on its own set of experimental groups $\theta_{\text{exp}}^\perp$.

This environment can be resolved completely to make predictions for 'real world' results. In this case, both the pre- and post-assignment ergodic distributions differ from $P_{-\infty}$ and must be re-computed. With an additional statement that describes feasible histories that accord with $Y^\perp$, the results of carrying out the social experiment in the novel environment also fall out automatically. In other words, a hypothetical experiment $\mathcal{E}^\perp$ can be designed and simulated on an hypothetical environment $\mathcal{V}^\perp$ based on information gathered from an experiment conducted in another environment.

When demographic variables are segregated from unobserved and estimated parameters as in done in the definition of a social environment, then out-of-population predictions are simpler, if not necessarily better or more accurate. Only the hypothetical policy vector must be observed. Without making *ad hoc* assumptions, the estimated exogenous vectors are inserted into the state space as the values pertaining to the hypothetical environment. The individual's problem is solved for the policy vector for each unobserved type $k$. What remains to be determined are the type proportions $\lambda^\perp$. This requires a vector $y^\perp$ of moments drawn from the experimental group with at least $K - 1$ elements. In the worst

31

case this vector could simply be the mean outcomes from the unselected population as long as the measurement vector is longer than the number of unobserved types. The type proportions are then chosen to match the data:

**Definition 6. Out-of-population Environment.** Given an estimated social environment $\hat{\mathcal{V}}$, an *out-of-population environment* $\hat{\mathcal{V}}^{\perp} \left( Y^{\perp}, y^{\perp}, \Psi_{p}^{\perp}, \mathcal{H}^{\perp} \left[ y^{\perp}; \theta_{\text{cond}} \right], \hat{\mathcal{V}} \right)$ is defined as:

$$\hat{\mathcal{V}}^{\perp} = \left( \hat{\Theta}^{\perp}, \mathbf{A}(\theta), U(\alpha, \theta), P \left\{ \theta' \, | \, \alpha \, , \, \theta \, \right\}, Y^{\perp}(\alpha, \theta), y^{\perp} \right)$$

where

$$\hat{\theta}^{\perp} \equiv \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & \theta_{\text{end}} & \hat{\theta}_{\text{exog}}^{\perp} & \Psi_{p}^{\perp} \end{pmatrix}$$

$$\hat{\theta}_{\text{exog}}^{\perp} = \begin{pmatrix} \hat{\Lambda}^{\perp} & \hat{\Gamma} \end{pmatrix}$$

$$\hat{\Lambda}^{\perp} = \arg \min_{\Lambda^{\perp} \in \Xi^{K}} \check{\Delta}(\Lambda^{\perp}) A^{\perp} \check{\Delta}(\Lambda^{\perp})' \tag{46}$$

$$\Psi_{x}^{\perp} = \begin{pmatrix} 0 & \mathcal{H}^{\perp} \left[ y^{\perp}; \theta_{\text{cond}} \right] & \cdot & \cdot & \cdot \end{pmatrix}.$$

The presence of $\cdot$ in a definition indicates parameters that are irrelevant to the definition. A out-of-population environment has no experimental treatment program, but the single group $e = g = G = E = 0$ must be defined to indicate how the data are generated. The length of the selection period $T$ is zero, which is a normalization since no exogenous variation needs to be inserted into the probability distribution as in (32). (The sample may still be selected after $t_0$ as indicated by $\mathcal{H}^{\perp}[\cdot]$.) The data-gathering process must be described because the out-of-population mixture across types is estimated using GMM on the available data.

Individual problems need to be solved exactly $K$ times in order to solve $\hat{\lambda}^{\perp}$ iteratively. The extent to which estimates from a social experiment can or should be applied in other contexts depends on the quality of data available from elsewhere and, of course, the ability of the estimated model to predict behavior for out-of-sample policies.

Segregating observed variation from unobserved variation and using a finite-mixture to model unobserved variation leads to straightforward applications of the model to populations not eligible for the experiment. In contrast, the more common way of including

32

policy and demographic information parameters within a model is to interact the in-sample index $d$ directly with the exogenous parameter vector $\Gamma$. For example, consider a case in which requirement R2 is violated by including a dummy-variable for, say, race in a wage equation. To apply the estimated model out of sample would require that race is observed in $Y^{\perp}$ or the effect of race must be integrated out of the predictions. In either case, it is assumed that race plays the same role in the target population as in the original population. This assumption is not necessary under R2, which would assume that people categorized as different races behave in observably different ways because they are a different mixture of the underlying types $k$.

### V.C Other Extensions

The appendix describes four other extensions: attrition bias, measurement error, non-stationary environments and stochastic policies, and equilibrium environments.

## VI. Conclusion

All economic policy interventions, whether carried out as an experiment or not, change the incentives and constraints faced by the agents in the economy. These changes create a backdrop of 'exogenous variation' against which to measure individual and aggregate responses. The literature on social and natural experiments usually casts the analysis of exogenous variation in a static framework, although all such interventions have dynamic effects, whether intended or not.

A commonly held preconception is that experiments and quasi-experiments can be interpreted without a model of subject response to the experiment. This position has been challenged by others on both practical and theoretical grounds. This paper folds these objections into a general framework that defines an experiment as an unexpected change in an environment that is already dynamic and uncertain. The element of surprise

is a maintained hypothesis about the distribution of unobserved states at the time of assignment to experimental groups. The environments and experiments in which the intuition from static frameworks carries through can be described. They are found to be a small and uninteresting subset of the set of environments and experiments defined here. As an alternative to these unappealing assumptions, a constructive set of tools is developed here for designing, estimating, and applying a complete model of a social experiment under general conditions.

To implement this framework for a particular experiment can entail a great deal of notation and computational costs. The return is a model that is parsimonious in free parameters. In particular, the model can be seen as a restriction placed on the analysis of experimental impact, which in itself has no power to predict outside the experiment under either the null hypothesis that the model is true or the alternative. In addition, the cost of extending the estimated model to other environments is relatively cheap. The solution and estimation of the social experiment is an upfront cost which makes multiple applications inexpensive on the margin. Whether the computational cost and attention to detail demanded by the framework are worthwhile depends on the quality of the situation under study and the model chosen to explain it, as well as how the results are to be put to use.

# VII. Appendix

## Notation

**N1.** Greek letters denote vectors and structural parameters of the model (elements of the exogenous vector). Vectors are often split into sub-vectors. For example, let $x_a = (x \quad y \quad z)$ and $x_b = (r \quad s)$. Then $x = (x_a \quad x_b)$ is equivalent to saying $x = (x \quad y \quad z \quad r \quad s)$. Capital Greek letters are typically a set or the concatenation of vectors with the corresponding lower-case letter. A concatenation of items can also be denoted $[x_n]$, where $x_n$ could be a scalar, a vector or a set.

**N2.** Lower-case Roman letters denote individual elements of vectors; capital Roman letters denote policy parameters or the maximum value of a variable with corresponding lower-case letter; bold capital letters denote sets.

**N3.** The cardinality of a set $\mathbf{S}$ is denoted $|\mathbf{S}|$, and the length of a vector $v$ is denoted $|v|$. The open simplex in $M$ dimensions is denoted

$$\Xi^M \equiv \left\{ x : x \in \Re^M,\ i < M \to 0 < x[i] < 1 - \sum_{j=1}^{i-1} x[j],\ x[M] = 1 - \sum_{i=1}^{M-1} x[i] \right\}. \tag{47}$$

**N4.** The Boolean (indicator) function is denoted

$$\mathcal{B}[z] \equiv \begin{cases} 0 & z \text{ is false} \\ 1 & z \text{ is true.} \end{cases} \tag{48}$$

**N5.** If $\theta$ is a vector, then $\theta\big|_{x=5,y=3}$ means that the elements of $\theta$ named $x$ and $y$ are to be set to $5$ and $3$, respectively. When the variable being set is clear from the value being assigned, its name is dropped: $\theta\big|_Y$ is short for $\theta\big|_{y=Y}$. Finally, $\theta\big|_{\neg x}$ is defined as the vector excluding the variable $x$.

**Transitions.** Let $h$ index variables in $\theta$ and let the discrete jump process for state variable $\theta[h]$ be written $\theta[h]^\star(\bar{h}, [\pi_j^h], [\mathbf{H}_j])$. Then

$$P\{\theta' \,|\, \alpha, \theta\} = \prod_{h=1}^{|\theta|} \left[ \left(1 - \sum_{j=1}^{J^h} \pi_j^h\right) \mathcal{B}\left[\theta'[h] = \bar{h}\right] + \sum_{j=1}^{J^h} \frac{\pi_j^h}{|\mathbf{H}_j|} \mathcal{B}\left[\theta'[h] \in \mathbf{H}_j\right] \right]. \tag{49}$$

Combining choice and transition probabilities generates the complete state-to-state transition

$$P_s\{\theta' \,|\, \theta\} = \sum_{\alpha \in \mathbf{A}(\theta)} P\{\alpha \,|\, \theta\} P\{\theta' \,|\, \alpha, \theta\} \tag{50}$$

$$= \sum_{\alpha \in \mathbf{A}(\theta)} \left\{ \frac{\tilde{v}(\alpha, \theta)}{\sum_{\alpha'} \tilde{v}(\alpha', \theta)} \prod_{h=1}^{|\theta|} \left[ \left(1 - \sum_{j=1}^{J^h} \pi_j^h\right) \mathcal{B}\left[\theta'[h] = \bar{h}\right] + \sum_{j=1}^{J^h} \frac{\pi_j^h}{|\mathbf{H}_j|} \mathcal{B}\left[\theta'[h] \in \mathbf{H}_j\right] \right] \right\}.$$

## Proof of I1.

**PI1a.** Since $\Theta$ and $A(\theta)$ are both finite sets, $U(\alpha, \theta)$ is bounded. By R1.1 unique values of $v(\alpha, \theta)$ and $V(\theta)$ solve the contraction (1)-(2). By R1.3 all feasible actions have non-zero choice probabilities at each state. Combined with R1.2 and R2.1 the state-to-state transition $P_s$ is 'irreducible' over all states in $\theta_{\text{end}}$ excluding $d$ and $k$. Thus a unique stationary distribution over these states exists (e.g. see Theorem 3.11.1 in Judd 1998). By R2.4-6 all other elements of the state vector are single-valued, so a unique $P_{-\infty}(\theta)$ exists for which (11) and (12) hold.

**PI1b.** From R2.3, R1.3, and envelope theorems on the dynamic program, $P\{\alpha|\theta\}$ and $P_s$ are smooth functions of $\theta_{\text{exog}}$. Since $P_{-\infty}$ uniquely solves a set of linear equations in $P_s$ it also is smooth in $\theta_{\text{exog}}$.

## Proof of I2.

**PI2a.** In a natural experiment $g$ is directly related to $d \in \theta_{\text{end}}$, and is thus not iid at $t_0$.

**PI2b.** Immediate.

## Proof of I3. Solution algorithm.

**a.** Set $d = D$, $k = K$, $f = F$, $r = 1$, $g = G$, $e = E$.

**b.** Set policy parameters $\Psi_d = \theta_{\text{pol}}[d]$.

**c.** Set exogenous parameters $\Gamma = \theta_{\text{end}}[k]$.

**d.** Iterate on (2) for all $\theta_{\text{end}}$ except $k$ and $d$ to find $V(\theta)$ to a specified tolerance.

**e.** Compute choice probabilities ($P\{\alpha|\theta\}$ in ) and solve the stationary distribution ($P_{-\infty}$ in ) for the current value of $k$ and $d$. See Judd (1998 p. 85) for details.

**f.** Decrease $f$ by 1 and set $r = R[f]$. Iterate back to $r = 1$ solving for value functions and choice probabilities.

**g.** Repeat the previous step through $f = 1$.

**h.** If $e = 0$, $g < G$ and $\tau > 0$, then set $V_0 = V_F$ iterate on the infinite horizon problem in () to a specified tolerance.

**i.** Set $t = t_{\min}$. Compute $E_{-\infty}[Y \mid d]$, $\omega(k; \theta_{\text{cond}})$ and $\Omega\{\theta|k, \theta_{\text{cond}}\}$.

**j.** Increase $t$ by 1. Update $\Omega$ and $\omega$. If $t = t_0$ then reset clock to $\bar{\theta}_{\text{clock}}$.

**k.** Repeat previous step until $t = t_{\max}$.

**l.** Decrease $g$ by 1. Set $f = F$ and return to step $f$ until $g = 0$.

**m.** Decrease $e$ by 1. Set $g = G - 1$ and return to step $f$ through $e = 1$.

**n.** Decrease $k$ by one. Set $e = E$ and return to step $c$ through $k = 1$.

**o.** Compute $E[Y] = \sum_k \lambda[d, k] E[Y]$ for all $g, e, t$.

**p.** Decrease $d$. Set $k = K$ and return to step $b$ through $d = 1$.

## Proof of I4.

**PI4a.** Since $d$ and $k$ are the only invariant states the remaining endogenous states form an ergodic system. The distribution over states conditional on $k$ and $d$

converges to $P_{-\infty}\{\theta \mid k, d\}$. Since $k \notin \theta_{\text{cond}}$, the conditional distribution must correct for the permanent selection bias in $k$.

**PI4b.** This follows from the previous result, because

$$\lim_{t \to \infty} \Delta(\theta_{\text{cond}}) = \sum_{k=1}^{K} \left( \lambda_0^\star(k|\theta_{\text{obs}}) - \lambda_0^\star\left(k|\theta_{\text{obs}}\Big|_G\right) \right) E_{-\infty}[Y \mid k, d] \tag{51}$$

which is not a function of the treatment program.

**PI4c.** If treatment is randomized $\lambda_0^\star(k|\theta_{\text{obs}}) = \lambda_0^\star\left(k|\theta_{\text{obs}}\Big|_G\right)$. If the environment is homogeneous $\lambda_0^\star(1|\theta_{\text{obs}}) = 1$. In either event the ultimate impact is zero.

**Proof of I5.** The question is when can treatment impact be expressed as a difference in results integrated over a common distribution: $P\left\{\theta \mid (\alpha^0, \theta^0)\right\} = P\left\{\theta \mid (\alpha^0, \theta^0)\right\}\Big|_G$. Because if this is the case we can write

$$\Delta_S(\alpha^0, \theta^0) = \sum_\theta \left[ E[Y \mid \theta] - E[Y \mid \theta]\Big|_G \right] P\left\{\theta \mid (\alpha^0, \theta^0)\right\}. \tag{52}$$

This would mean that the impact of the treatment is isolated in conditional choice probabilities at the time of measurement. To show the two conditions are each sufficient, express the conditional distribution recursively:

$$\tilde{P}\left\{\theta^1 \mid (\alpha^0, \theta^0)\right\} \equiv P^\star\left\{\theta^1 \mid (\alpha^0, \theta^0)\right\}$$
$$\text{for } s > 0,$$
$$\tilde{P}\left\{\theta^{s+1} \mid (\alpha^0, \theta^0)\right\} \equiv \sum_{\theta^s} P_s\left\{\theta^{s+1} \mid \theta^s\right\} \tilde{P}\left\{\theta^s \mid (\alpha^0, \theta^0)\right\}$$

**PI5a.** Assignment at $t_0$ takes place after $\alpha^0$ is set, so for $s = t_0 + 1$ there is no distributional impact, except for the difference in the initial clock settings. That is, $P\left\{\theta^1 \mid (\alpha^0, \theta^0)\right\} = P\left\{\theta^1\Big|_{\theta_{\text{clock}}} \mid (\alpha^0, \theta^0)\Big|_{g'}\right\}$ for any other group $g'$.

**PI5b.** For $t > t_0 + 1$ there will be an impact on choice probabilities during periods $t_0 + 2, t_0 + 3, \ldots, t - 1$ unless behavior is irrational. Otherwise, only when the environment and the treatment are autonomous will the intervening impact have no effect on transition probabilities: $P_s\{\theta^{s+1}|(\alpha^0, \theta^0)\} = P_s\left\{\theta^{s+1}|(\alpha^0, \theta^0)\Big|_{g'}\right\}$.

**Proof of I6.**

**I6a.** Immediate, because the mean and all individual treatment effects are zero.

**I6b.** This is equivalent to I5, but selection on permanent unobserved heterogeneity must be ruled out by assuming homogeneity. Otherwise the mean treatment effect is averaging over unobserved types while an individual conditions on their own type.

**I6c.** In the convoluted scenario, the chain of conditional probabilities is exactly that defined in (32) and (34). One gets the same numerical value when a person is randomly drawn from the stationary distribution, conditional on $d$, and forms the correct conditional expectation. However, this fails to maintain the information set available to the individual, since the randomly selected person knows their current state (including their type $k$).

### Coordinating time across groups.

**a.** For $e > 0$, let the initial clock have elements $(t_0 + 1 \quad r_c \quad f_c)$. Let $\bar{r}(e) \equiv r_c + \sum_{x=1}^{f_c - 1} R(x)$ be an index of how far into the program the group enters treatment. Set $t_0(z) \equiv 0$, where $z \equiv \arg\max_{\theta_{\exp}} \bar{r}(\theta_{\exp})$ is the group that enters treatment at the latest stage.

**b.** Set $t_0$ in other groups so that members of the group cannot reach treatment state $\bar{\theta}_{\text{clock}}(z)$ before $t = 0$: $t_0(\theta_{\exp}) \equiv -\left[\bar{r}\left(\bar{\theta}_{\text{clock}}(z)\right) - \bar{r}\left(\theta_{\exp}\right)\right]$.

**c.** Set $t_{\min}$ to be $1 + t_0 - T(e)$. Set $t_{\max}$ as $t_0$ plus the length of post-assignment observations in the group.

**d.** For $e = 0$, set $f' = f^\star(1,0,0)$ for $(\alpha, \theta) \in h_{\bar{t}}$, and $\bar{t} = \min_e t_{\min}(e) - 1$.

### Applications of the Framework.

#### Ap1. The Illinois Re-employment Bonus Experiment.

Woodbury and Spiegelman (1987) report the results of two controlled experiments carried out on new claimants to unemployment insurance (UI) in Illinois. Meyer's (1995) review of the evidence from these and related experiments suggests an on-going influence of such experiments and the analysis of their results. The two experiments were based on the same entry conditions and shared a common control group, thus $G = 3$. In each case the treatment consisted of a cash bonus for finding a job within 11 weeks of beginning their claim. In one treatment group the bonus was paid to the employee. In the other it was paid to the employer. It was indeed an unexpected experimental intervention with a finite duration, thus $E > 0$. There were not multiple samples with different entry conditions, thus $E = 1$. The main selection criterion was that the subject filed a new UI claim and registered in a job service area in northern and central Illinois. The claimant also needed to qualify for 26 or more weeks of UI and to be between the ages of 20 and 55. For simplicity we will assume that location and age are treated as demographic variables in order to concentrate on the unemployment criteria. Let $m, j \in Y(\alpha, \theta)$ be indicators for employment status and registering at a job service bureau in the current period (week), respectively. Let $n \in Y(\alpha, \theta)$ be the number of weeks of UI eligibility if unemployed next period. The modeler decides whether and to what extent these variables are under the control of the subject by locating them in $\theta$ or $\alpha$ or by making them functions of other elements of the outcome. The length of the selection period is thus $T = 2$: the date of random

assignment is $t_0 = 0$ and selection begins at $t_{\min} = -1$. The selection criteria are, for $t = -1$, $\mathcal{H}[y; \theta_{\text{cond}}] = m(n > 26)$ and, for $t = 0$, $\mathcal{H}[y; \theta_{\text{cond}}] = j(1 - m)$. There are three phases of treatment ($F = 4$). Phase $f = 1$ is the period during which the subject can qualify for the employment bonus. They must start a job within 11 weeks of entry into the experiment, while receiving UI benefits for the spell of unemployment at the start of the experiment. This implies $R[1] = 11$. The phase increment is $f_+ = m$ until $r = 11$, when it becomes $f+ = 2(1 - m)$. That is, the person stays in the qualification phase as long as they don't start working. If they reach the final period they move to phase $4$ if they don't work, otherwise they make the default transition to $f = 2$. The second phase ($f = 2$) requires the subject to hold the job for four periods, thus $R[2] = 4$. The phase increment is $f_+ = 2(1 - m)$ for $r < 4$ because if they failed to keep the job for four weeks they lost the chance to receive the bonus and moved back to the real world. At $r = 4$ the phase increment is $f_+ = 1 - m$. Phases 1 and 2 are qualification phases because the utility and transition are not treated during them. The last treatment phase lasts one period, $R[3] = 1$, and is characterized by an increase of \$500 in income, paid either directly to the individual (and presumably entering $U(\alpha, \theta)$) or to the employer, depending on the value of $g$.

By mapping the re-employment bonus experiment into this framework it becomes possible to analyze its outcome using a model of endogenous job search and UI eligibility. That is, the re-employment bonus will affect the propensity to enter UI-eligible unemployment is embedded within the model. The experiment provides no direct experimental (i.e. controlled) evidence on how the bonus would affect this propensity, because both the treatment and control groups entered this state under the status quo. However, a model cast in this framework would need to explain both the treatment and control group outcomes while accounting for endogenous transitions into unemployment. If successful at crossing this in-sample hurdle, it is not a great leap for the model to predict accurately the unobserved situation of a permanent application of the treatment. The finite and temporary nature of the bonus, and the fact that its receipt depends on subject action,is fully accounted for by using the pre-assignment value of outcomes in the post-treatment phase $F$.

We can contrast this approach to the one taken by Davidson and Woodbury (1993) that lacked a framework to embed an experiment in a social environment. In their analysis, becoming unemployed and becoming eligible for UI are exogenous parameters calibrated to pre-assignment outcomes and results from other studies. The analysis thus assumes a homogeneous population and an unselected group entering the experiment. Policy implications are then based on an assumption that implementing a reemployment bonus would leave unchanged the propensity to become unemployed and to establish UI eligibility prior to becoming unemployed. While the Davidson and

Woodbury model is a stationary infinite horizon environment, experimental outcomes are analyzed as if the re-employment bonus immediately became an on-going (recurring) policy.

**Ap2. NIT Experiments**

Unlike the Illinois re-employment bonus experiments, the NIT experiments were not cleanly designed and executed. Here we will just describe how many of the problems with the NIT experiments listed by Pencavel (1987) are normalized by this framework. First comes the strongly selected and widely scattered geographic locations of the experiments. This is automatically captured by allowing each demographic group to have its own distribution over unobserved type. (Applying results to non-sampled locations is described in section VII.) Next, is the fact that entry into the experiment was based on being a low income household. This simply requires that $T(e) = 1$, income be an element $Y(\alpha, \theta)$, and the feasible condition be that income is less than some amount. The three-year period of treatment made it unclear whether short-run or long-run elasticities were being measured. In the social experiment the finite nature of the treatment is explicit, the difference with a permanent policy innovation is clear, and the permanence of the treatment effects would be endogenous to the sequential decision model and the parameters estimated from the data. Next, in some locales the size of the transfer and the claw-back rate differed across treatment households. This is handled by specifying more than one treatment group ($G > 1$). Finally, in some locales assignment into groups was not random, because households were assignment treatment according to forecasts designed to minimize the experimental costs. As long as these non-random aspects of assignment are explicit, the jump process $g_0^\star()$ can be specified to handle them. In short, the framework outlined above coupled with a adequate model of labor supply decisions can rescue the exogenous variation generated by the NIT experiments from the so-called flaws in its design.

**Ap3. Vietnam Draft**

A classic example of a natural experiment is the Vietnam draft lottery (Angrist 1990). The experiment is that each draft board drew birth dates randomly to determine priority in filling their draft requirement. Hence, assignment to treatment was based on a demographic variable. The draft lottery was not a true experiment, because the structure of the lottery was known ahead of time and the 'experiment' could be anticipated by the subjects. Thus a description of the draft lottery as a social experiment requires $E = 0$. However, treatment was temporary because exposure to the draft was not permanent, and thus $\tau = 0$.

The draft lottery is naturally randomized if date of birth is independent of unobserved type $k$. Note that other demographic variables, such as age and parent's education, can be included in the social environment and are not required to be independent of birthdate. Yet, a model of the Vietnam draft

lottery as a naturally randomized social experiment resting on this assumption is problematic on two levels. First, policies such as minimum school entry age, mandatory school attendance age, and minimum automobile license age, are directly related to birthdate. They may therefore generate a correlation between birthdate and other endogenous variables. However, the social experiment is not bound to account for these policy differences, and it may be acceptable to ignore them.

At a deeper level the unobserved type $k$ may be correlated with birthdate. For instance, suppose the model includes unobserved human capital and observed schooling as endogenous choices. Schooling at age 16 can be treated as a demographic variable, but differences in human capital at age 16, conditional on schooling at age 16, would in this environment be treated as exogenous parameters that differ over unobserved type $k$. The modeler has the choice of assuming that the distribution over $k$ is independent of birthdate, which would result in the draft lottery being a naturally randomized experiment. The fact that local draft boards conducted their own lotteries can be exploited to avoid assuming that $k$ is independent of birthdate. A birthdate in one locale with a low draft number could be paired with the same birthdate in another locale with a high draft number. Location and date of birth are perhaps safely assumed to be uncorrelated, and it may be possible to pair locales with similar policy vectors (e.g. similar schooling policies). The result is a social environment in which the draft is a naturally randomized experiment as defined earlier.

**Ap4. Non-experimental Lifecycle Data**

The requirement that the social environment is stationary is not restrictive in non-experimental situations. Consider the case: $E = 0$; $\tau = 0$; $U(\alpha, \theta)\big|_{f=0} = U(\alpha, \theta)\big|_F = 0$. The 'experiment' can now represent an individual's lifetime. Iteration on the infinite horizon value function converges in two iterations. Different phases correspond to different stages of decision making generated by policy, such as the age before majority and the age after mandatory retirement. Transitions must be specified so that the desired initial distribution over endogenous states corresponds to the ergodic distribution $P_{-\infty}\{\theta\}$. Otherwise, assuming a transparent environment means that the initial distribution is not relevant to applying the model to data, because initial conditions are observed up to the unobserved type $k$.

**Other Extensions of the Basic Framework.**

**Ex1. Attrition Bias.** In practice, individuals often agree to participate in an experiment, but then when their assignment is revealed they make a second choice to continue or not. When participation in the experiment is costly to the subjects, and the treatment is valued differently than the status quo, this leads to a bias in estimated impacts. Related to this issue is attrition bias: if

subjects drop from study at a rate that depends on the treatment, then the attrition rate will differ between the treatment and control group.

It is straightforward to allow for endogenous attrition in the social experiment. Let $\alpha$ contain a choice variable $j$ which equals one when the subject agrees to participate in the next period. Thus, for $t \geq t_0$, $\mathcal{H}[y; \theta_{\mathrm{cond}}] = j$. Typically $j = 0$ implies that $f' = F$. That is, non-participation means treatment ends and the subject returns to the real world. Suppose further that the cost of participation is an additive component of utility:

$$u^\star(\alpha, \theta) = j\left(u(\alpha, \theta) - \kappa\right) + (1 - j)u(\alpha, \theta)\Big|_{f=F} \tag{53}$$

where $\kappa \in \Upsilon$ and $u^\star(\alpha, \theta)$ is the utility augmented by the participation decision. Given that the subject has participated up to the current period, they will participate again if the current value plus discounted future expected value of treatment outweighs the cost $\kappa$. Only when agents are myopic does the participation decision boil down to a static tradeoff: $j = 1$ when $u(\alpha, \theta) - u(\alpha, \theta)\Big|_{f=F} > \kappa$. In any case, the decision depends on all the exogenous parameters and the realized state a subject finds themselves in. For example, subjects with less to gain from the treatment (perhaps because they are less patient and have lower values of $\delta$) are more likely to drop out conditional on the realized state.

To control for attrition that occurs at the point of random assignment can require a slight extension of the model, because this bias is caused by a lag between a subject's agreeing to participate in the experiment and being assigned treatment status $g$. To account for this lag, $t_0$ now represents the period at which the experiment is revealed and participation is requested. Phase $f = 1$ would now denote the period between the initial decision to participate and the date of random assignment. $R[1]$ is the maximum lag between a baseline survey and random assignment. Phase $2$ is the first true phase of treatment. During phase $1$, $g$ is a discrete jump process that will take on a new value when the subject makes the jump out of phase 1. Phase 1 would in most cases by an entry period as defined above. The extended utility $u^\star$ implies that the decision to drop out (and, hence, random assignment bias) is a function of the discounted expected value of treatment and the status quo situation all conditional upon information available to the subject. Thus, subjects who initially agreed to participate based on (53) may drop out once their group assignment is realized. (If full records are kept at $t_0$, often called the baseline interview, then the initial decision to participate can also be modelled using $u^\star$.)

**Ex2. Measurement Error.** As discussed earlier, maximum likelihood estimation is sensitive to the assumption that measurements are made without error. True measurement error and discrepancies between reality and the model can generate zero-probability observations that leave the likelihood function

undefined. This can be seen in the main component of the likelihood defined in the Appendix which contains the Boolean component $\mathcal{B}\left[y_n(\theta_{\text{cond}}) = Y(\alpha, \theta)\right]$. We can extend the model by specifying that outcomes are subject to independent normally-distributed errors. The exogenous vector is expanded to take the form:

$$\theta_{\text{exog}} \equiv \left( \Lambda \quad \Gamma \quad \Sigma \right) \tag{54}$$

where $\Sigma$ is a $M$-vector of standard deviations in results, $\Sigma = [\sigma_m]$. The Boolean component is now replaced with a normal kernel over the outcome space:

$$\exp\left\{ -\frac{1}{2} \left( \frac{y_n\left(\theta_{\text{cond}}\right) - Y(\alpha, \theta)}{\Sigma} \right) \left( \frac{y_n\left(\theta_{\text{cond}}\right) - Y(\alpha, \theta)}{\Sigma} \right)' \right\}. \tag{55}$$

(The division by $\Sigma$ is element-by-element, and a scaling constant would also appear.) The measurement error would not affect the progress of the experiment through the transition rule $f_+(y; \theta_{\text{clock}})$. The likelihood function is not reliant on 'transparency' of the state and action vector, so the same general form applies in the case of this type of measurement error.

**Ex3. Non-stationary Environments and Stochastic Policies.** If at some level the environment is not ergodic then it is not clear how experimental results, however obtained, could be useful *ex post* to address policy questions. When subject behavior is non-stationary before and after the experiment, then the same subjects will never be in the same position to react to true policy changes that occur after the experiments end. A good reason to relax the assumption of stationarity is to allow subjects to make decisions while facing a finite horizon, for example in anticipation of retirement or death. Another good reason is to have individuals make decisions starting from an initial distribution over states that is not endogenous to the model. As discussed earlier, lifecycle data that does not include random assignment can be modelled as a program of treatment within a special (trivial) stationary environment. But to include lifecycle effects and experimental treatment, an endogenous variable, say $a$, must be cyclical: $a' = a^\star(\mathcal{B}\left[a < A\right] a + 1, [0], [\emptyset])$. The value of terminal states, $V(\theta\big|_A)$ and the jump values of initial states, $P(\theta\big|_{a=1})$, are both determined by exogenous parameters and not by endogenous choices. Dying agents are 're-born' at $a = 1$ and the overall stationarity required for drawing lessons for policy is restored through an overlapping generations framework. The allowance for policy innovations as $e = 0$ means that the transition path between the old and new steady-states are computable within the framework already described.

Recently Keane and Wolpin (2002) have examined the effect of uncertain changes in government policies within a model of welfare. In the basic definition of a social environment the policy vector takes on a singe value, and the demographic variable $d$ serves as a permanent index into the policy vector. It is straightforward (albeit expensive) to allow the values in the policy vector

to follow a discrete jump process, $d' = d^{\star}(d, [\pi_d], [\{d\}])$. The possible realized values of policy parameters would be pre-determined and the parameters of the discrete jump process could be estimated inside or outside the model. As with Keane and Wolpin's framework, this allows expectations that heretofore unobserved regimes (such as introduction or elimination of a program) can affect current behavior. Policy experiments would then be conducted by letting the government set the realized values of policy parameters and/or their transition probabilities.

**Ex4. Equilibrium Environments.** Equilibrium can be considered a restriction on the environment that subjects face. For example, in a partial equilibrium environment, prices are parametric to all individual agents and the modeler. In a general equilibrium the prices remain parametric to agents but become endogenous to the modeler's choice or estimates of underlying preference and technology parameters. In many cases we can associate equilibrium outcomes with certain elements of the policy vector $\Psi_p$. For example, rather than treat prices as given, they would have to satisfy an endogenous restriction: $h(\Psi_p) = 0$, where all the implied optimizing behavior and aggregate outcomes are implicit in $h(\cdot)$. Solving the model would require finding exogenous parameters that satisfy this restriction. Alternatively, the econometric objective can be penalized for a failure to satisfy these restrictions. In this approach the equilibrium prices are elements of the exogenous vector. Ferrall (2004) discusses the computational tradeoffs between these two approaches to imposing equilibrium in a general class of problems of which a social experiment is a special case.

Combining equilibrium restrictions and small-scale experiments require special consideration (Heckman et al. 1998). The framework here makes it straightforward to include equilibrium responses only when appropriate. For $f = F$ one equilibrium restriction is imposed within and perhaps across demographic groups. Any case of $e > 0$ is an experiment within the same equilibrium, presuming $N(\theta_{\text{obs}})$ is small. A case of $e = 0$ will require computing new equilibrium restrictions whether it is an hypothetical or real policy innovation being analyzed.

# VIII. References

Aguirregabiria, V. and P. Mira. 2002. "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models," *Econometrica* 70, 4, 1519 - 1543.

Angrist, J. D. 1990. "Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records," *American Economic Review* 80, 3, June, 313-336.

Arcidiacono, P. and J. B. Jones 2003. "Finite Mixture Distribution, Sequential Likelihood, and the EM Algorithm," *Econometrica* 71, 3, 933 - 946.

Burtless, G. "The Case for Randomized Field Trials in Economic and Policy Research," *Journal of Economic Perspectives* 9, Spring 2, 1995.

Davidson, C. and S. A. Woodbury 1993. "The Displacement Effect of Reemployment Bonus Programs," *Journal of Labor Economics* 11, 4, 575-605.

Eckstein, Z. and Wolpin, K. I. (1990a), "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey," *Journal of Human Resources*, XXIV, 562-598.

Eckstein, Zvi and Wolpin, Kenneth I. 1999. "Why Youths Drop out of High School: The Impact of Preferences, Opportunities, and Abilities," *Econometrica 67*, 6, pp. 1295-1339.

Ferrall, Christopher. 2004. "Solving Finite Mixture Models in Parallel," manuscript, Queen's University, `http://qed.econ.queensu.ca/pub/faculty/ferrall/papers/mixtures.pdf`.

Heckman, J. J. and Singer, B. 1984. "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data," *Econometrica* 52

Heckman, J. J. and J. A. Smith, "Assessing the Case for Social Experiments," *Journal of Economic Perspectives* 9, 2, 1995.

Heckman, J. J., R. J. Lalonde and J. A. Smith 1999. "The Economics and Econometrics of

Active Labor Market Programs," in *Handbook of Labor Economics*, 3a, O. Ashenfelter and D. Card (eds.), Elsevier;

Heckman, J. J., L. Lochner, and C. Taber. 1998. "General Equilibrium Treatment Effects: A Study of Tuition Policy," *American Economic Review* 88, 2, 381-386.

Heckman, J. J., H. Ichimura, and Petra Todd 1998. "Matching as an Econometric Evaluation Estimator," *Review of Economic Studies* 65, 2, 261-294.

Hotz, V. Joseph, Guido W. Imbens and Julie H. Mortimer 1999. "Predicting the Efficacy of Future Training Programs Using Past Experiences," NBER working paper T0238.

Ichimura, H and C. Taber 2000. "Direct Estimation of Policy Impacts," IFS working paper 0005, `http://www.ifs.org.uk/workingpapers/wp0005.pdf`.

Imbens, Guido W. and Joshua D. Angrist 1994. "Identification and Estimation of Local Average Treatment Effects," *Econometrica* 62, 2, 467-475.

Judd, Kenneth L. 1998. *Numerical Methods In Economics*, Cambridge, Mit Press.

Keane, M. P. and K. I. Wolpin. 2001. "Estimating Welfare Effects Consistent with Forward-Looking Behavior. Part I: Lessons from a Simulation Exercise, PIER working paper 01-019, University of Pennsylvania.

Lise, Jeremy, Shannon Seitz, and Jeffrey Smith. 2004. "Equilibrium Policy Experiments and the Evaluation of Social Programs," NBER Working Paper No. w10283, `http://-papers.nber.org/papers/W10283`.

Manski, Charles F. 1996. "Learning about Treatment Effects from Experiments with Random Assignment of Treatments," *Journal of Human Resources* 43, 4, 709-733.

Meyer, B. D. 1995. "Lessons from the U.S. Unemployment Insurance Experiments," *Journal of Economic Literature* 33, 1, pp. 91-131.

Moffitt, Robert A. 2004. "The Role of Randomized Field Trials in Social Science Research," *American Behavioral Scientist*, 47, 5, 506-540.

Orcutt, G. H. and A. G. Orcutt 1968. "Incentive and Disincentive Experimentation for Income Maintenance Policy Purposes," *American Economic Review* 58, 4, 754-772.

Pencavel, John 1987. "Labor Supply of Men: A Survey," in *Handbook of Labor Economics*, Vol. 1, Orley C. Ashenfelter and Richard Layard (eds.), North-Holland.

Rosenzweig, M. R. and K. I. Wolpin 2000. "Natural 'Natural Experiments' in Economics," *Journal of Economic Literature* 38, 4, 827-874.

Rust, J. 1994. "Structural Estimation of Markov Decision Processes, in *Handbook of Econometrics*, Volume 4, R. Engle and D. McFadden (eds.), 30823139, North Holland.

—— 1997. "Using Randomization to Break the Curse of Dimensionality," *Econometrica* 65, 3, 487-516.

Shearer, B. 2004. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment," *Review of Economic Studies* 71, 2, 513-534.

Todd, Petra and Kenneth I. Wolpin. 2003. "Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility: Assessing the Impact of a School Subsidy Program in Mexico," PIER Working Paper No. 03-022, `http://-www.econ.upenn.edu/Centers/pier/Archive/03-022.pdf`

Woodbury, S. A. and R. G. Spiegelman 1987. Bonuses to Workers and Employers to Reduce Unemployment: Randomized Trials in Illinois," *American Economic Review* 77, 4, 513-30.

'