

Estimation and Inference in Social Experiments

Christopher Ferrall

Department of Economics
Queen's University

<http://qed.econ.queensu.ca/pub/faculty/ferrall>

August 26, 2002

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.

Keywords: Treatment Effects, Impact Analysis, Dynamic Programming, Policy Experiments

Research support from the Social Science and Humanities Research Council of Canada is gratefully acknowledge. This paper is partly based on a draft research report supported by the Social Research and Demonstration Corporation. Helpful comments from Shannon Seitz, Bruce Shearer, and seminar participants at the 2001 CEA Meetings in Montreal, Laval, the Chicago Fed, and the 2002 SED Meetings in New York are appreciated.

I. Introduction

The Negative Income Tax (NIT) experiments of the 1960s and 1970s economists 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. By the late 1980s economists had become less ambivalent about experiments, in part because the goals stated for social experiments had become more modest. Social experiments are now more carefully designed to demonstrate the impact of a particular policy change on selected outcomes (Meyer 1995 and Heckman et. al 1999). 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. The framework developed here addresses several problems that plagued evaluations of the NIT and later experiments by building a model of how forward-looking agents become eligible for an experiment and how they react to being randomly assigned to receive treatment.¹ This forward-looking model is compared to the static framework in which experiments are typically discussed and analyzed. Several implications are presented that suggest the intuition from the static model does not readily extend into the dynamic model. In short, while random assignment to treatment alone is useful it is not sufficient on

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 (2001) experiment on incentives to carry out hard physical labor (tree planting).

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 value unless combined with a model that accounts for how subjects become eligible for random assignment and why they respond to treatment. [Hotz et. al \(1999\)](#) find that results of one experiment are not good predictors of results of related but not identical experiments. [Rosenzweig and Wolpin \(2000\)](#) and [Ichimura and Taber \(2000\)](#) scrutinize the randomization hypothesis underlying instrumental variables when used to make policy-related inferences. Ichimura and Taber focus on a static simultaneous equations framework, which has limited capacity to account for information available to subjects as they make decisions. Rosenzweig and Wolpin use simple models to illustrate problems in interpreting results based from natural experiments.

What is new in this paper is a unified framework for studying reactions to unanticipated events. The formal definition of a social experiment developed here 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. For a complete application of the framework to analyze an experiment, see [Ferrall \(2002b\)](#).

The model of individual behavior is a discrete-choice dynamic program, and as such this paper relates closely to [Rust \(1994\)](#) which characterized the estimation of discrete-choice Markov decision problems. Here three state variables are filled into the standard model to account for experimental (or quasi-experimental) treatment. The standard model is also extended to integrate unanticipated 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.

As with any fully specified ‘structural’ model, a complete likelihood function for a sample of independent observations is generated, and some results concerning identification of the estimated parameters are provided. However, particular attention is also paid to estimating the model using averaged data. There are three reasons. First, maximum likelihood estimation on individual data demands the model be completely consistent with the experiment and data collection. Rectifying even incidental discrepancies between the data and the data generating process often requires costly increases in computation. Applying Generalized Method of Moments (GMM) on averaged data smooths many discrepancies between the model and reality, allowing a more parsimonious model. Second, use of averaged data means the estimated model is nested by a atheoretic impact analysis. The social experiment model becomes a null hypothesis about mean outcomes which, if not rejected, can replace the atheoretic analysis as an explanation of the experimental results. Of course, maximum likelihood is itself a GMM estimator, so emphasizing GMM based on limited information is motivated by a perception that structural estimation is a non-nested alternative to a simple analysis of impact. Here it is emphasized that auxiliary assumptions can be viewed as testable against an alternative that simply treats random assignment as exogenous variation. Third, many controlled and natural experiments are studied with proprietary and confidential data. In these cases access to individual outcomes is limited. This paper emphasizes that limited access to outcomes in evaluations of experiments does not in itself preclude estimation of a rational model that can explain and predict experimental outcomes.

II. Experiments and Econometrics

II.A Background

Heckman et al (1999) and Rosenzweig and Wolpin (2000) provide comprehensive surveys of experimental and quasi-experimental studies in economics. Analysis begins by selecting a result to study, denoted Y . The impact of treatment on Y is the difference Δ between Y conditional on assignment to the treatment group ($g = 0$) and assignment to the control group ($g = 1$):²

$$\Delta \equiv Y_0 - Y_1. \tag{1}$$

Δ is treated as an unobserved random variable endogenous to many decisions made before and after the experiment begins. In the extreme, when only the random assignment to experimental groups is considered exogenous, the sample mean treatment effect,

$$\hat{\Delta} = \hat{E}[Y_0 \mid g = 0] - \hat{E}[Y_1 \mid g = 1] \tag{2}$$

remains an unbiased estimate of the population mean individual treatment effect:

$$E[\hat{\Delta}] = E[Y_0 \mid g = 0] - E[Y_1 \mid g = 1] = E[Y_0] - E[Y_1] = E[\Delta]. \tag{3}$$

Thus $\hat{\Delta}$ answers an important question. Yet a well-run experiment raises more questions than it answers. How would Δ change if the ‘dose’ were increased or decreased? Or if subjects were treated longer? How long after the end of the experiment do its effects last? When Δ is taken as a primitive parameter, these questions are unanswerable. They are answerable, however, within a model of behavior of the subjects in both the treatment and control groups. The remainder of this section argues that this approach is valuable for several other reasons.

² Usually the subscript 1 denotes the treatment group, but using it to denote the control group reflects more clearly the algorithm to solve the model.

II.B Too Many Impacts

One advantage of using only mean treatment effects from controlled experiments is that they are easy to explain to politicians and other non-specialists (Burtless 1995). In the textbook case, Y and Δ are both scalars, and explaining the impact of a program would be straightforward. In practice, several measurements are taken at different points after random assignment. If $E[\Delta]$ is estimated for all of them, then Y is a vector that concatenates all results of interest measured at each point in time. With a multi-dimensional impact vector, no one statistic summarizes the outcome of the experiment. Because Δ is not based on a theory of how subjects react to the experiment, impact analysis provides no objective way to aggregate variation across time or across impacts. The clarity of a scalar Δ gives way to a cloud of impacts out of which many different and possibly conflicting conclusions might be reached.

Although the number of measurements in an experiment is usually pre-determined, presumably the impact continues past this point. The reported impact vector Δ itself does not predict or forecast further results. Put succinctly, impact analysis is a method-of-moments estimate with a free parameter for every result of interest (Δ). This freedom leaves no power to forecast behavior after the experiment or subsequent to (hypothetical) policy changes. Correlations across impacts cannot be exploited to generate more efficient estimates. In the model developed here all results arise from one underlying model with a number of parameters that are independent of the measurements take from the experiment. Impact within the experiment can be projected beyond the measured sample into hypothetical situations, including policies not identical to the treatment and populations not included in the experiment.

Another rhetorical advantage of using only Δ is timeliness. An impact analysis can be available almost as soon as the data are collected. In contrast, estimating a dynamic model using data on individuals is notoriously time-consuming and sensitive to assumptions. However, the advantages of social experiments apply equally well, if not more so, to a model-

based analysis. The exogenous variation generated by the experiment compresses the length of time and the size and diversity of the sample required to achieve a given precision in parameter estimates.

II.C Sample-Selection and Heterogeneity Bias

Experiments are rarely conducted on randomly sampled members of the population of interest. It is more cost-effective to sample subjects who are most likely to respond to the treatment. This limits the extent to which results of an experiment apply to the unselected population. In addition, while the mean treatment effect is non-parametrically identified from a controlled experiment, it is unlikely that the mean summarizes the implications for policy (Heckman and Smith 1995). The problems of sample-selection and heterogeneous impacts are two aspects of the more fundamental problem of heterogeneity in unobserved parameters, which include preferences and lagged decisions. If, in a given experiment, these issues seem exceptionally large, they can be avoided by narrowing the question of interest to the effect of treatment on the treated.

In contrast, a model of the experiment (if it is good) applies in both the treatment and control groups, both inside and outside the experiment. From the point of view of a subject, Δ is a random variable that depends on: the state the subject was in at the point of random assignment, realized events between random assignment and the point of measurement, and the subject's optimal response to the experimental treatment. In short, permanent and transitory heterogeneity in treatment effects is also endogenous to the model even before introducing parameter heterogeneity across subjects. As long as the selection criteria for the experiment are clear and relatively simple, correcting for sample selection becomes an implication of the model.

II.D Experiments and Policy

The fundamental question posed by an experiment is:

If the treatment were to replace the status quo and become policy, would the experimental results predict the real results?

This is different than the more narrow ‘what if’ questions raised early, because they could, in principle, be answered by a larger or longer experiment that varied parameters. But any social experiment must be finite-lived, voluntary, and without an open window of opportunity to enter treatment. A policy, on the other hand, applies in the future to subjects not eligible at the time the policy was announced. It may be available repeatedly, and it may not be voluntary. Because experimental treatment is a one-shot temporary intervention, the subjects of an experiment all anticipate a return to a world without treatment. The status quo policy is *de facto* the terminal state for both the treatment and control groups.

From the beginning this limitation was noted, and [Orcutt and Orcutt \(1968\)](#) suggested that varying the treatment period may identify the size of the terminal bias in treatment. However, if the current policy were eliminated, the sample selected into an experiment comes from a world which would no longer exist. The initial conditions themselves cannot be treated away. Without a reliable model of behavior under the status quo it is not possible to correct for sample selection bias in real policy forecasts solely based on the experimental results. But if the experimental treatment were implemented, it would modify the status quo and by that fact alone modify the impact of the program. When combined with the other issues left unresolved by random assignment alone, status quo bias in both the initial and terminal conditions of an experiment make it nearly impossible to project experimental results into policy predictions without a model of behavior to explain reaction to the treatment.

III. 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, poli-

cies 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.

III.A Notation

Greek letters denote vectors and structural parameters of the model (elements of the exogenous vector). Capital Greek letters are typically 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. 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.

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

$$\Xi^M \equiv \left\{ x : x \in \mathbb{R}^M, i < M \rightarrow 0 < x[i] < 1 - \sum_{j=1}^{i-1} x[j], x[M] = 1 - \sum_{i=1}^{M-1} x[i] \right\}. \quad (4)$$

The Boolean (indicator) function is denoted

$$\mathcal{B}[z] \equiv \begin{cases} 0 & z \text{ is false} \\ 1 & z \text{ is true.} \end{cases} \quad (5)$$

If θ is a vector, then $\theta|_{x=5,y=3}$ means that the elements of θ 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: $\theta|_Y$ is short for $\theta|_{y=Y}$. Finally, $\theta|_{-x}$ is defined as the vector excluding the variable x .

III.B Subject Behavior

The model of subject behavior is a discrete-state, discrete-choice dynamic program.³

Each period the subject is in a state θ , an element of the state space Θ . Each period

³ 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, [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.

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

$$v(\alpha, \theta) = U(\alpha, \theta) + \delta E[V(\theta')] = U(\alpha, \theta) + \delta \sum_{\theta'} P\{\theta' | \alpha, \theta\} V(\theta'), \quad (6)$$

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

$$\forall \theta \in \Theta, \quad V(\theta) = \max_{\alpha \in \mathbf{A}(\theta)} v(\alpha, \theta). \quad (7)$$

The dynamic program in (6) and (7) 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 θ as the concatenation of five sub-vectors, each of which contains state variables that play specific roles in the model:

$$\begin{aligned} \theta &\equiv \left(\theta_{\text{clock}} \quad \theta_{\text{exp}} \quad \theta_{\text{end}} \quad \theta_{\text{exog}} \quad \theta_{\text{pol}} \right) \\ &= \left((t \quad r \quad f) \quad (g \quad e) \quad (\dots \quad d \quad k) \quad (\Lambda \quad \Gamma) \quad [\Psi_p[d]] \right). \end{aligned} \quad (8)$$

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 (6)-(7). 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 describe the subject’s situation inside the experiment, and they are defined in the next section. Outside the experiment the values of θ_{clock} and θ_{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 $P\{\theta' | \alpha, \theta\}$ 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** $q' = q^*(\bar{q}, [\pi_j], [\mathbf{Q}_j])$ means $\forall (\alpha, \theta)$,

$$P\{q' | \alpha, \theta\} = \sum_{j=1}^{J(\alpha, \theta)} \pi_j(\alpha, \theta) \frac{\mathcal{B}[q' \in \mathbf{Q}_j(\alpha, \theta)]}{|\mathbf{Q}_j(\alpha, \theta)|} + \left(1 - \sum_j \pi_j(\alpha, \theta)\right) \mathcal{B}[q' = \bar{q}(\alpha, \theta)]. \quad (9)$$

In words, $q^*(\bar{q}, [\pi_j], [\mathbf{Q}_j])$ 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 π_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 (9) makes it explicit that the arguments of the jump process depend on the current outcome (α, θ) .

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

- ◇ *absorbing at n* if $q = n \rightarrow q' = q^*(n, [0], [\emptyset])$;
- ◇ *invariant* if q is absorbing for all values $q = 1, 2, \dots, Q$;
- ◇ *autonomous [at θ]* if $[\text{at } \theta] q^*$ is not a function of α ;
- ◇ *independent [at $q = n$]* if $(q = n \text{ implies}) q^*$ is not a function of α or any other state variable $s \neq q$;
- ◇ *ergodic [at θ]* if for $n = 1, 2, \dots, Q$, $\exists \alpha^n \in \mathbf{A}(\theta) : P\{q' = n | (\alpha^n, \theta)\} > 0$;
- ◇ *iid* if for constants $[p_i]$, $q' = q^*\left(q, \left[p_1 \quad \dots \quad 1 - \sum_{i=1}^{Q-1} p_i\right], [\{1\} \quad \dots \quad \{Q\}]\right)$;
- ◇ *dependent* if q is not ergodic, but for any two values q^*, q^{**} , there exists a chain of outcomes $(\alpha^1, \theta^1), (\alpha^2, \theta^2), \dots, (\alpha^m, \theta^m)$, $m < Q-1$, with corresponding values of q

such that $P\{q' = q^{**} \mid \cdot\} > 0$, $P\{q' = q_m \mid (\alpha^{m-1}, \theta^{m-1})\} > 0, \dots, P\{q' = q_1 \mid (\alpha^*, \theta^*)\} > 0$;

To show the flexibility of this notation, consider one further 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 as

$$q^* \left(q, \mathcal{B}[z = 0] \left[p_1 \quad p_2 \quad \dots \quad 1 - \sum_i^{Q-1} p_i \right], [\{1\} \quad \{2\} \quad \dots \quad \{Q\}] \right). \quad (10)$$

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.

Invariant states as defined above differ across individuals but are fixed for a given individual. The set of invariant states partition 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

$\Upsilon \in \mathbf{U} \subset \mathbb{R}^{N_1}$; $P\{\theta' \mid \alpha, \theta\}$ is a smooth function of a vector $\Pi \in \mathbf{P} \subset \mathbb{R}^{N_2}$ of exogenous parameters; \mathbf{U} and \mathbf{P} are bounded open sets.

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 pa-

parameter $0 < \rho < 1$:

$$\begin{aligned} \tilde{v}(\alpha, \theta) &\equiv \mathcal{B}[\alpha \in \mathbf{A}(\theta)] \exp \left\{ \frac{\rho}{1-\rho} [v(\alpha, \theta) - V(\theta)] \right\} \\ P\{\alpha | \theta\} &= \frac{\tilde{v}(\alpha, \theta)}{\sum_{\alpha'} \tilde{v}(\alpha', \theta)}. \end{aligned} \quad (11)$$

From R1 the complete transition function can be described. Let h index variables in θ and let the discrete jump process for state variable $\theta[h]$ be written $\theta[h]^*(\bar{h}, [\pi_j^h], [\mathbf{H}_j])$. Then

$$P\{\theta' | \alpha, \theta\} = \prod_{h=1}^{|\theta|} \left[\left(1 - \sum_{j=1}^{J^h} \pi_j^h \right) \mathcal{B}[\theta'[h] = \bar{h}] + \sum_{j=1}^{J^h} \frac{\pi_j^h}{|\mathbf{H}_j|} \mathcal{B}[\theta'[h] \in \mathbf{H}_j] \right]. \quad (12)$$

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

$$\begin{aligned} P_s\{\theta' | \theta\} &= \sum_{\alpha \in \mathbf{A}(\theta)} P\{\alpha | \theta\} P\{\theta' | \alpha, \theta\} \\ &= \sum_{\alpha \in \mathbf{A}(\theta)} \left\{ \frac{\tilde{v}(\alpha, \theta)}{\sum_{\alpha'} \tilde{v}(\alpha', \theta)} \prod_{h=1}^{|\theta|} \left[\left(1 - \sum_{j=1}^{J^h} \pi_j^h \right) \mathcal{B}[\theta'[h] = \bar{h}] + \sum_{j=1}^{J^h} \frac{\pi_j^h}{|\mathbf{H}_j|} \mathcal{B}[\theta'[h] \in \mathbf{H}_j] \right] \right\}. \end{aligned} \quad (13)$$

III.C 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 α 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 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. θ_{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 (6)-(7) are collected in a vector of size $N = N_1 + N_2 + 2$ that is specific to each type k :

$$\Gamma[k] \equiv (\Upsilon_k \quad \Pi_k \quad \delta_k \quad \rho_k). \tag{14}$$

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

$$\theta_{\text{pol}} \equiv (\Psi_p[1] \quad \Psi_p[2] \quad \dots \quad \Psi_p[D]). \tag{15}$$

- R2d.** The (invariant) proportion of the population with a given combination (k, d) equals $\bar{\lambda}[d]\lambda[d, k]$, where: $\bar{\lambda}[d] \in \Psi_p[d]$; $\lambda[d, k] \in \theta_{\text{exog}}$; $\bar{\lambda} \in \Xi^D$; and $\lambda[d] \in \Xi^K$ for all d .
- R2e.** Outside an experiment θ_{clock} and θ_{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:

$$\begin{aligned}
\theta_{\text{exog}} &\equiv (\lambda[1, 1] \ \cdots \ \Lambda[D, K] \ \Upsilon_1 \ \Pi_1 \ \delta_1 \ \rho_1 \ \cdots \ \Upsilon_K \ \Pi_K \ \delta_K \ \rho_K) \\
&\equiv (\Lambda[1] \ \Lambda[2] \ \cdots \ \Lambda[D] \ \Gamma[1] \ \cdots \ \Gamma[K]) \\
&\equiv (\Lambda \ \Gamma).
\end{aligned} \tag{16}$$

The bounded open set of possible exogenous vectors is:

$$\hat{\Theta}_{\text{exog}} \equiv \{\Xi^K\}^D \times \left\{ \mathbf{U} \times \mathbf{P} \times (0, 1)^2 \right\}^K \subset \mathfrak{R}^P. \tag{17}$$

The exogenous vector contains all parameters of the model to be estimated from data. The $\hat{\cdot}$ is included in $\hat{\Theta}_{\text{exog}}$ to avoid confusing the set of valid estimates of θ_{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 Γ and proportions Λ 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 θ_{pol} without changing $\bar{\Lambda}$ or Λ 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 θ_{pol} .

The primitive outcome-to-state transition and the endogenous choice probabilities combine to generate the state-to-state transition P_s in (13). 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 Λ . 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

$$\forall d, k, \quad \bar{\lambda}[d]\lambda[d, k] = \sum_{\theta'} \mathcal{B}[k = k', d = d'] P_{-\infty}\{\theta'\} \quad (18)$$

$$\forall \theta \in \Theta, \quad P_{-\infty}\{\theta\} = \sum_{\theta'} P_s\{\theta | \theta'\} P_{-\infty}\{\theta'\}. \quad (19)$$

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

Proof.

P1a. Since Θ 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 (6)-(7). 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 θ_{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 (18) and (19) hold.

P1b. From R2.3, R1.3, and envelope theorems on the dynamic program, $P\{\alpha | \theta\}$ and P_s are smooth functions of θ_{exog} . Since $P_{-\infty}$ uniquely solves a set of linear equations in P_s it also is smooth in θ_{exog} .

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.

III.D Measurements

From a subject's point of view the outcome (α, θ) 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. $Y : A(\Theta) \times \Theta \rightarrow \mathfrak{R}^M$ is the measured result vector generated by outcomes. $Y(\alpha, \theta)$ does not include k or θ_{exog} .

R3b. The state variables in the vectors θ_{obs} and θ_{cond} , defined as

$$\theta_{\text{obs}} \equiv (\theta_{\text{exp}} \quad d) = (g \quad e \quad d) \tag{20}$$

$$\theta_{\text{cond}} \equiv (t \quad \theta_{\text{obs}}) = (t \quad g \quad e \quad d), \tag{21}$$

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 θ_{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 θ is

$$E[Y | \theta] \equiv \sum_{\alpha \in A(\theta)} P\{\alpha | \theta\} Y(\alpha, \theta). \quad (22)$$

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

$$E_{-\infty}[Y | d] \equiv \sum_{\theta'_{\text{end}}} \mathcal{B}[d = d'] E[Y | \theta'] P_{-\infty}\{\theta'\} / \bar{\lambda}[d] \quad (23)$$

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 θ_{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

$$\nu(\Theta, \mathbf{A}(\theta), U(\alpha, \theta), P\{\theta' | \alpha, \theta\}, Y(\alpha, \theta), y)$$

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

Examples. A social environment ν is

- ◇ *autonomous* if no subject can affect the future: $\forall(\alpha, \theta), P\{\theta' | \alpha, \theta\} = P\{\theta' | \theta\}$;
- ◇ *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 $\theta_{\text{end}}|_{-k,d} \in Y(\alpha, \theta)$.

In a myopic environments 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 θ_{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. \(1998\)](#) 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.

IV. The Experiment

IV.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 θ_{clock} in (8). 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 Ψ_t .

Requirement 4. Program of Treatment.

R4a. The treatment vector takes the form:

$$\Psi_t = (R \quad f_+(y; \theta_{\text{clock}}) \quad \dots) \quad (24)$$

where R is a $F - 1$ vector of positive integers denoting the maximum length of phases $f = 1, 2, \dots, F - 1$, 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 Ψ_t , which can interact with the current phase f and treatment group g .

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

$$f' = f^*(f + \mathcal{B}[r = R[f]] + f_+, 0, \emptyset) \quad (25)$$

$$r' = r^*(\mathcal{B}[f_+ = 0, r < R[f]] r + 1, 0, \emptyset) \quad (26)$$

$$t' = t^*(t + 1, 0, \emptyset). \quad (27)$$

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

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

$$U(\alpha, \theta) = U(\alpha, \theta) \Big|_{\tilde{r}, \tilde{t}} \quad (28)$$

$$q^*(\bar{q}, [\pi_j], [\mathbf{Q}_j]) = q^*(\bar{q}, [\pi_j], [\mathbf{Q}_j]) \Big|_{\tilde{r}, \tilde{t}}. \quad (29)$$

The transition rules for treatment are simpler than the notation might suggest. The phase next period can only depend on θ_{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, $y_n(\theta_{\text{cond}})$.

Examples of treatment.

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

- ◇ a *qualification phase* if $\forall (\alpha, \theta), U(\alpha, \theta) = U(\alpha, \theta)|_F$ and $\forall q \in \theta_{\text{end}}, q^*() = q^*()|_F$;
- ◇ *autonomous* if $f_+ = 0$;
- ◇ a *waiting period* if it is an autonomous qualification phase;
- ◇ an *entry period* if it is a waiting period and $f = 1$;
- ◇ *voluntary* if $\forall \theta \exists \alpha \in A(\theta) : f + \mathcal{B}[r = R[f]] + f_+ = F$.

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.

IV.B Experimental Groups and Policy Innovations

Each subject is associated with an experimental status described by the sub-vector θ_{exp} introduced in (8). 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 $\bar{\theta}_{\text{clock}}(\theta_{\text{exp}})$.

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 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, θ_{cond} defined in (21). Valid measurements are indicated with a Boolean function $\mathcal{H}[y; \theta_{\text{cond}}]$. 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 $\mathcal{H}[y; \theta_{\text{cond}}] = 1 - \mathcal{B}[z]$. 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, $g = g^*(g, [\mu_g], [g])$. Recall that the arguments of g^* can depend upon the current outcome (α, θ) . 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 μ_g is under the control of the experimenter, which implies that μ_g can only be a function of d , e , and the result vector $Y(\alpha, \theta)$. A ‘randomized experiment’ μ_g only depends on e and would thus equal the sample proportions of each treatment group. 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 α . 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 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 stationarity 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 Ψ_x ,

$$\Psi_x(\theta_{\text{exp}}) = \left(T \quad \mathcal{H}[y; \theta_{\text{obs}}] \quad g_0^*(\cdot) \quad \bar{\theta}_{\text{clock}} \quad \tau \right), \quad (30)$$

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

R5b. For $g = G$, $\bar{\theta}_{\text{clock}} = (t_0 + 1 \quad 1 \quad F)$. Based on $T(e)$ and other values of $\bar{\theta}_{\text{clock}}$, observations are split into pre- and post-assignment periods, $t_{\min}(e) \leq t_0(e) \leq 0 \leq t_{\max}(e)$, 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 :

$$E > 0, t = \bar{t} \rightarrow e' = e^*(e, 1/E, [e]) \quad (31)$$

$$t \neq t_{\min} \rightarrow e' = e^*(e, 0, \emptyset) \quad (32)$$

$$t \neq t_0 \rightarrow g' = g^*(g, 0, \emptyset) \quad (33)$$

$$t = t_0 \rightarrow g' = g_0^*(\cdot). \quad (34)$$

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

$$\begin{aligned} t < t_{\min}(e) &\rightarrow \mathcal{H}[y; \theta_{\text{cond}}] = 1 \\ t \leq t_0(e) &\rightarrow \mathcal{H}[y; \theta_{\text{cond}}] = \mathcal{H}[y; \theta_{\text{cond}}] \Big|_{g'} \end{aligned} \quad (35)$$

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

$$f = F \rightarrow f' = f^*(F, \mathcal{B}[e = 0, g < G] \tau(g), 0). \quad (36)$$

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

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 θ_{cond} such $\mathcal{H}[y; \theta_{\text{cond}}] \neq 1$;
- ◇ *randomized* if at $t = t_0$, g^* 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 : $\forall g \exists \mathbf{D}_g \subseteq \{0, 1, \dots, D\} : t = t_0, d \in \mathbf{D}_{\tilde{g}} \rightarrow g' = g^*(\tilde{g}, 0, \emptyset)$;
- ◇ *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 θ_{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 $\theta_{\text{end}}|_{-d,k}$), 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.

Proof.

P5a. In a natural experiment g is directly related to $d \in \theta_{\text{end}}$, and is thus not iid at t_0 .

P5b. Immediate.

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 unantici-

pated 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.

IV.C Examples of Social Experiments

Before considering the properties of impact analysis and other econometric techniques applied to a social experiment as defined here, the practical use of the framework is illustrated by mapping four situations into it.

IV.C.1 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 θ or α 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.

IV.C.2 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 $g_0^*(\cdot)$ can be specified to handle them. In short, the framework outlined above coupled with an adequate model of labor supply decisions can rescue the exogenous variation generated by the NIT experiments from the so-called flaws in its design.

IV.C.3 Vietnam Draft

A classic example of a natural experiment is the Vietnam draft lottery ([Angrist 1990](#)). The experiment is that each draft board drew birthdates 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.

IV.C.4 Non-experimental Lifecycle Data

The requirement that the social environment is stationary is not restrictive in non-experimental situations. Consider the case: $E = 0$; $\tau = 0$; $U(\alpha, \theta)|_{f=0} = U(\alpha, \theta)|_F = 0$. 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 .

V. Predicting Experimental Results

Return to section II, where $\hat{E}[Y_0 \mid g = 0]$ and $\hat{E}[Y_1 \mid g = 1]$ denoted the sample mean outcomes within treatment groups. Impact was defined in (1) as the difference in mean outcomes. 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 θ_{cond} introduced in (21). 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.

V.A Sample Selection

The model of subject behavior generates a stationary distribution across states outside the experiment, denoted $P_{-\infty}\{\theta\}$ and defined in (19). 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

$$\Omega\{\theta' \mid k, \theta_{\text{cond}}\} \equiv \frac{1}{\max\{E, 1\} \lambda[d] \lambda[k, d]} \sum_{\theta} P_s\{\theta' \mid \theta\} \mathcal{B}[k] P_{-\infty}\{\theta\} \Big|_{\bar{t}}. \quad (37)$$

Requirement R5.2 and the presence of E (the number of experimental groups) in the de-

nominator 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 θ_{cond} . For all k this proportion is initialized as

$$\omega(k; \theta_{\text{cond}}) \Big|_{\bar{t}} \equiv 1. \quad (38)$$

For computational purposes, $\omega(k; \theta_{\text{cond}})$ is *not* defined as the distribution of k given θ_{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 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

$$P^* \{ \theta' | \theta \} = \sum_{\alpha} P \left\{ \theta' \Big|_{\mathcal{B}[t=t_0] \bar{\theta}_{\text{clock}}} \mid \alpha, \theta \right\} \mathcal{H}[Y(\alpha, \theta); \theta_{\text{cond}}] P \{ \alpha | \theta \}. \quad (39)$$

The assignment $\mathcal{B}[t = t_0] \bar{\theta}_{\text{clock}}(\theta_{\text{exp}})$ is shorthand for saying that the insertion of the initial clock setting for θ' 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 θ at the start of any period $t \geq t_{\text{min}}$ is denoted $\Omega \{ \theta | k, \theta_{\text{cond}} \}$. Subjects will choose actions that determine their outcome

(α, θ) 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 $\mathcal{H}[y; \theta_{\text{cond}}] = 0$. The cumulative proportion of the type k population making it to t is

$$\omega(k; \theta_{\text{cond}}) = \omega\left(k; \theta_{\text{cond}} \Big|_{t-1}\right) \left[\sum_{\theta'} \sum_{\theta} P^* \{ \theta' | \theta \} \Omega \left\{ \theta | k, \theta_{\text{cond}} \Big|_{t-1} \right\} \right]. \quad (40)$$

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

$$\Omega \{ \theta' | k, \theta_{\text{cond}} \} = \frac{\omega\left(k; \theta_{\text{cond}} \Big|_{t-1}\right)}{\omega(k; \theta_{\text{cond}})} \sum_{\theta} P^* \{ \theta' | \theta \} \Omega \left\{ \theta | k, \theta_{\text{cond}} \Big|_{t-1} \right\}. \quad (41)$$

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:

$$\lambda^*(k | \theta_{\text{cond}}) \equiv \lambda[k, d] \frac{\omega(k; \theta_{\text{cond}})}{\sum_{k'=1}^K \omega(k'; \theta_{\text{cond}})}. \quad (42)$$

At $t_{\min} - 1$ 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 θ_{exp} . If selection ends at t_0 , then λ^* becomes constant for all $t > t_0$:

$$t > t_0 \rightarrow \lambda^*(k | \theta_{\text{cond}}) = \lambda^*\left(k | \theta_{\text{cond}} \Big|_{t_0+1}\right) \equiv \lambda_0^*(k | \theta_{\text{obs}}). \quad (43)$$

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.

V.B Mean Outcomes and Mean Treatment Effects

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

$$\hat{E}[Y | \theta_{\text{cond}}] = \frac{1}{N(\theta_{\text{cond}})} \sum_{n=1}^{N(\theta_{\text{cond}})} \mathcal{B}[y_n(\theta_{\text{cond}}) \neq \dot{y}] y_n(\theta_{\text{cond}}) \quad (44)$$

$$\hat{\Delta}(\theta_{\text{cond}}) \equiv \hat{E}[Y | \theta_{\text{cond}}] - \hat{E}[Y | \theta_{\text{cond}}|_G], \quad (45)$$

where $N(\theta_{\text{cond}})$ was defined in R3. Using (41), the model analogue to the empirical values are

$$E[Y | \theta_{\text{cond}}] = \sum_{\theta} \lambda^*(k|\theta_{\text{cond}}) \Omega\{\theta | k, \theta_{\text{cond}}\} E[Y|\theta] \quad (46)$$

$$\Delta(\theta_{\text{cond}}) \equiv E[Y | \theta_{\text{cond}}] - E[Y | \theta_{\text{cond}}|_G]. \quad (47)$$

The impact in (47) 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 $E_{-\infty}[Y | d]$ defined in (23) 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 $t \rightarrow \infty$. 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 θ_{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 $\tau = 0$,

$$\forall \theta, \quad P_{+\infty}\{\theta \mid \theta_{\text{cond}}\} \equiv \lim_{t \rightarrow \infty} \lambda_0^*(k|\theta_{\text{obs}}) \Omega\{\theta \mid k, \theta_{\text{cond}}\} = \left(\frac{\lambda_0^*(k|\theta_{\text{obs}})}{\lambda[k, d]} \right) P_{-\infty}\{\theta \mid d\}. \quad (48)$$

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

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

Proof.

P5a. 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 .

P5b. This follows from the previous result, because

$$\lim_{t \rightarrow \infty} \Delta(\theta_{\text{cond}}) = \sum_{k=1}^K \left(\lambda_0^*(k|\theta_{\text{obs}}) - \lambda_0^*(k|\theta_{\text{obs}}|_G) \right) E_{-\infty}[Y \mid k, d] \quad (49)$$

which is not a function of the treatment program.

P5c. If treatment is randomized $\lambda_0^*(k|\theta_{\text{obs}}) = \lambda_0^*(k|\theta_{\text{obs}}|_G)$. If the environment is homogeneous $\lambda_0^*(1|\theta_{\text{obs}}) = 1$. In either event the ultimate 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 perturbation 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

$$E_{+\infty}[Y \mid d, g] \equiv E \left[Y \mid \theta_{\text{cond}} \Big|_{e=f=0, g} \right]. \quad (50)$$

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

$$\Delta_{\infty}[d, g] \equiv E_{+\infty}[Y \mid d, g] - E_{-\infty}[Y \mid d]. \quad (51)$$

The notation and the substance behind (50) resolve both of the problems in using $P_{+\infty}$ to compute long-run impact. First, $E_{+\infty}[Y \mid d, g]$ 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.

V.C Individual Treatment Effects

Reaching farther back into section II we find the atheoretic outcomes Y_0 and Y_1 . Unlike the mean outcomes just discussed, these are two different outcomes for a single subject assigned to treatment and control status. The fundamental problem that social experiments attempt to solve is that Y_0 and Y_1 are never both experienced by the same subject. Yet, as with mean impacts, defining the model equivalents of Y_0 and Y_1 is not straightforward in a social environment. As a first pass consider:

$$Y_g \stackrel{?}{\equiv} Y(\alpha, \theta) \Big|_g. \quad (52)$$

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 (α, θ) in (52) 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 (52) 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 :

$$\Delta_S(\alpha^0, \theta^0) \equiv E \left[Y(\alpha, \theta) \mid (\alpha^0, \theta^0) \right] - E \left[Y(\alpha, \theta) \Big|_G \mid (\alpha^0, \theta^0) \right] \quad (53)$$

Notice that in the atheoretic framework $\Delta = Y_0 - Y_1$. It would be tempting to consider associating Δ_S with $E[Y(\alpha, \theta) - Y(\alpha, \theta) \big|_g]$. Random assignment would appear to imply that conditioning on assignment as of time t_0 has no effect on expectations, making it possible to interchange the expectation and subtraction operators in (53). 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 Ψ_t is autonomous **and** type k is autonomous or irrational,

then

$$\Delta_S(\alpha^0, \theta^0) = E \left[Y(\alpha, \theta) - Y(\alpha, \theta) \big|_G \mid (\alpha^0, \theta^0) \right].$$

Proof. 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

$$\Delta_S(\alpha^0, \theta^0) = \sum_{\theta} \left[E[Y \mid \theta] - E[Y \mid \theta] \big|_G \right] P \left\{ \theta \mid (\alpha^0, \theta^0) \right\}. \quad (54)$$

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:

$$\tilde{P} \left\{ \theta^1 \mid (\alpha^0, \theta^0) \right\} \equiv P^* \left\{ \theta^1 \mid (\alpha^0, \theta^0) \right\}$$

for $s > 0$,

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

P5a. Assignment at t_0 takes place after α^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,

$$P \left\{ \theta^1 \mid (\alpha^0, \theta^0) \right\} = P \left\{ \theta^1 \big|_{\theta_{\text{clock}}} \mid (\alpha^0, \theta^0) \big|_{g'} \right\} \text{ for any other group } g'.$$

P5b. For $t > t_0 + 1$ there will be an impact on choice probabilities during periods $t_0 + 2, t_0 + 3, \dots, t - 1$ unless behavior is irrational. Otherwise, only when the environment and the treatment are autonomous will the intervening impact have no effect on transition probabilities: $P_s\{\theta^{s+1} | (\alpha^0, \theta^0)\} = P_s\left\{\theta^{s+1} | (\alpha^0, \theta^0) \Big|_{g'}\right\}$.

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 θ_{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 Δ in (1) with Δ_S in (53).

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 (45) 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 empirical impact.

Implication 6: Subjective treatment effects. The mean treatment effect (45) corresponds to an expected individual effect (53) 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 $P_{-\infty}\{\theta | d\}$; 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 .

Proof

- 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 (39) and (41). 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).

What happened to the straightforward and intuitively appealing individual impact Δ as defined in (1). The problem is that (1) conditions 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.

VI. Estimation and Inference

VI.A Likelihood

The exogenous vector θ_{exog} introduced in (16) contains all parameters to estimate from the experimental results, including type proportions within demographic groups.

Implication 7: Likelihood. Under R1 - R5:

I7a. The post-assignment likelihood function for subject n of measurement group θ_{obs} is denoted $\mathcal{L}(\hat{\theta}_{\text{exog}}; n, \theta_{\text{obs}})$ and can be computed through *backward* recursion:

$$\mathcal{L}(\hat{\theta}_{\text{exog}}; \theta, t_{\text{max}} + 1, \theta_{\text{obs}}) \equiv 1$$

for $t_0 \leq t \leq t_{\text{max}}$,

$$\mathcal{L}(\hat{\theta}_{\text{exog}}; \theta, t, n, \theta_{\text{obs}}) = \sum_{\alpha} \left[\mathcal{B}[y_n(\theta_{\text{cond}}) = Y(\alpha, \theta)] \sum_{\theta'} P^* \{ \theta' \mid (\alpha, \theta) \} \mathcal{L}(\hat{\theta}_{\text{exog}}; \theta', t + 1, n, \theta_{\text{obs}}) \right]$$

$$\mathcal{L}(\hat{\theta}_{\text{exog}}; n, \theta_{\text{obs}}) \equiv \sum_{\theta'} \lambda_0^*(k', \theta_{\text{obs}}) \Omega \left\{ \theta' \mid k', \theta_{\text{cond}} \Big|_{t_0-1} \right\} \mathcal{L}(\hat{\theta}_{\text{exog}}; \theta', t_0 + 1, n). \quad (55)$$

I7b. The full-sample log-likelihood function is

$$\ln \mathcal{L}(\hat{\theta}_{\text{exog}}; y) = \sum_{\theta_{\text{obs}}} \sum_{n=1}^{N(\theta_{\text{obs}})} \ln \mathcal{L}(\hat{\theta}_{\text{exog}}; n, \theta_{\text{obs}}). \quad (56)$$

I7c. $\ln \mathcal{L}(\hat{\theta}_{\text{exog}}; y)$ is a smooth function of $\hat{\theta}_{\text{exog}}$ on $\hat{\Theta}_{\text{exog}}$.

I7d. In a transparent environment the likelihood simplifies to

$$\ln \mathcal{L} = \sum_{\theta_{\text{obs}}} \sum_{n=1}^{N(\theta_{\text{obs}})} \ln \left[\sum_{k=1}^K \lambda_0^*(k, \theta_{\text{obs}}) \prod_{t=t_0}^{t_{\text{max}}} P^* \{ \theta_{t+1} \mid (\alpha_t^n, \theta_t^n) \} P \{ \alpha_t^n \mid \theta_t^n \} \right], \quad (57)$$

where (α_t^n, θ_t^n) is shorthand for the unique solution to $Y(\alpha_t^n, \theta_t^n) = y_n(\theta_{\text{cond}})$.

I7e. In a transparent and homogeneous environment where Υ is identified, consistent estimates of Υ can be computed without solving the value function.

- I7f.** The likelihood for pre-random assignment results requires the joint distribution of θ for $t_{\min} \leq t < t_0$, the Cartesian product of the outcome space (excluding θ_{exp} and θ_{clock}) on itself $t_0 - t_{\min}$ times.
- I7g.** For an irrational type ($\rho_k \approx 0$) the discount factor δ_k and exogenous parameters Υ_k are statistically unidentified.

Proof

- I7a.** Backward induction
- I7b.** Immediate.
- I7c.** By Implication I1 all probabilities are smooth in θ_{exog} . $\Omega\{\theta|k, \theta_{\text{cond}}\}$ is smooth in $P\{\alpha|\theta\}$ and P_s , and the likelihood is smooth in Ω .
- I7d.** With transparency the state is observable up to k , and the sums over α and θ' in (57) collapse to a single factor.
- I7e.** With homogeneity Υ is identified from observable conditional choice probabilities.
- I7f.** By R1, $P\{\alpha|\theta\} = 1/|\mathbf{A}(\theta)|$ for an irrational type, which is not a function of Υ_k and δ_k . The likelihood would be flat in these parameters.

The ergodic distribution $P_{-\infty}$ accounts for the distribution of subjects across all invariant states, including unobserved type k which must be integrated out of the likelihood. The post-random assignment likelihood () integrates out pre-assignment results, because in general it is not feasible to exploit information prior to t_0 . The measurements at t_0 contribute to the likelihood even though assignment takes place after the measurement is made. The expression in () assumes that those losing eligibility at t_0 are not included in the sample. Although the distribution in (39) is relatively simple to compute for $t < t_0$, it only expresses the distribution across states for subjects *partially* eligible for their experimental group. The distribution across states in period $t < t_0$ for subjects who are *ultimately* eligible for random assignment requires tracking from t_0 back to period t . This goes against the Markov (sequential) nature of the subject's behavior and the selection criteria. The extra information

required to track the likelihood back from t_0 builds up exponentially with the length of the selection criteria. Unless the selection period is short, or the model itself has a small state space, the extended state space required by pre-assignment results will be infeasible to compute.

Rust (1994) provides the concentrated likelihood in transparent and homogeneous environments. Several other algorithms for estimating dynamic programming models assume a transparent and/or homogeneous environment. For example, Aguirregabiria and Mira (2002) maintain transparency and homogeneity. In a related paper, Arcidiacono and Jones (2002) relax homogeneity. Keane and Wolpin (1997) assume transparency but not homogeneity.

Because the utility and discount factor are unidentified under the null hypothesis of an irrational environment, inference is non-standard. However, some parameters in Π may be identified in an irrational environment. In particular, any state variables that are included in the result vector will allow identification of parameters of the discrete jump process that drives those variables.

VI.B Impact-based Estimates

There are several reasons why maximum likelihood estimates of the exogenous parameters may be infeasible or unattractive. First, the burden of calculating the likelihood may be large, particularly if the environment is transparent. When the environment is not transparent the subject's choices are based on more information than available in y . The likelihood must integrate out these choices. Second, the likelihood is exacting. Any inconsistencies between the model experiment and the actual experiment can lead to zero-probability events appearing in y . Finally, data on individual subjects may not be available to the researcher because experimental results are often proprietary due to informed-consent laws.

These concerns lead to consideration of estimating θ_{exog} from averaged data. Define $\check{\Delta}$ as the difference between predicted and observed mean outcomes:

$$\check{\Delta}(\theta_{\text{cond}}) \equiv \hat{E}[Y \mid \theta_{\text{cond}}] - E[Y \mid \theta_{\text{cond}}]. \tag{58}$$

By adding and subtracting a term, the empirical impact can be expressed as

$$\hat{\Delta}(\theta_{\text{cond}}) \equiv \check{\Delta}(\theta_{\text{cond}}) - \check{\Delta}(\theta_{\text{cond}}|_G) + \Delta(\theta_{\text{cond}}). \quad (59)$$

The empirical impact is a combination of theoretical outcomes and differences between theoretical and empirical outcomes. A set of reported impacts can be interpreted as reporting a particular mixture of theoretical and empirical moments. In an unrestricted impact analysis the mean outcomes in control groups are taken as given and not determined by $\hat{\Delta}(\theta_{\text{cond}})$. The impact for non-control groups is a parameter free available to match the means exactly. Let

$$N_M = \sum_{d=1}^D \sum_{e=0}^E \sum_{g=0}^G \sum_{t=t_0+1}^{t_{\max}} MB[N(\theta_{\text{cond}}) > 0] \quad (60)$$

denote the number of moments. The number of free parameters in an impact analysis is the number of impacts $N_U = N_M(G-1)/G$. A social experiment \mathcal{V} model places two kinds of restrictions on an impact analysis. First, a social experiment is a model of subject behavior where the number of parameters equals $P < N_U$. Second, the social experiment predicts the outcomes in control groups as well as treatment groups. This provides additional within-sample restrictions on the social experiment model.

Definition 6. Impact-based estimates of the exogenous vector solve

$$\hat{\theta}_{\text{exog}}^{\text{IE}} \equiv \min_{\theta_{\text{exog}} \in \hat{\Theta}_{\text{exog}}} \sum_{\theta_{\text{cond}}} \check{\Delta}(\theta_{\text{cond}})' A(\theta_{\text{cond}}) \check{\Delta}(\theta_{\text{cond}}). \quad (61)$$

where $A(\theta_{\text{cond}})$ is a $M \times M$ positive-definite matrix.

Implication 8: Impact-based estimation. Given an social experiment \mathcal{E} ,

- I8a.** The empirical impact in (45) can be considered an estimate of reduced-form parameters of the model, $\Delta(\theta_{\text{cond}})$.
- I8b.** $\hat{\theta}_{\text{exog}}^{\text{IE}}$ is a hypothesis nested by an impact analysis $\hat{\Delta}(\theta_{\text{cond}})$. The number of over-identifying restrictions on $\hat{\theta}_{\text{exog}}^{\text{IE}}$ is $\text{DF} \geq N_M - P - D$.
- I8c.** In an experiment with an entry period, the null hypothesis of a myopic environment implies the joint hypothesis $\Delta(\theta_{\text{cond}}) = 0$ for $t = t_0 + 1, \dots, t_0 + R(1)$.

Proof.

P8a. Immediate.

P8b. Immediate.

P8c. An entry period by definition has no affect on utility. In a static environment control and treatment groups will behave the same during the entry period, leading to a prediction of zero impact during the first $R(1)$ periods of the experiment.

In short, impact analysis cannot be rejected in favor of a social experiment model, \mathcal{E} , but it can fail to improve the fit enough to justify rejection of an estimated model using conventional tests. The final result also suggests that social experiment framework itself helps with identification of parameters of a dynamic program. Rust (1994) shows that a dynamic programming model is unidentified non-parametrically. The discount factor is a key parameter that itself cannot be identified without assuming parametric structure. The social experiment makes assumptions about the utility and transition, but by definition the presence of, say, an entry period is not a freely estimated parameter. Thus, within this structure P8c indicates that the value of the discount factor can yield testable implications about the outcomes of experiments.

VII. Applications and Extensions

VII.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. I8 says that an impact analysis uses up all degrees of freedom in the information contained in mean outcomes concerning the social environment, leaving no freedom to predict outcomes beyond the end of the sample, $t > t_{\max}$. By contrast, predictions using the social experiment framework are based on estimates of the exogenous

parameters, for example based on $\hat{\theta}_{\text{exog}} = \hat{\theta}_{\text{exog}}^{\text{IE}}$. For any $t > t_0$,

$$\hat{E}_t [Y] = \sum_{\hat{\theta}} \hat{E} [Y \mid \theta] \hat{\Omega}\{\hat{\theta} \mid k, \theta_{\text{cond}}\}. \quad (62)$$

Further, point estimates of the exogenous parameters come with an estimated distribution, $\hat{H}(\hat{\theta}_{\text{exog}}^{\text{IE}})$. Subject to onerous calculations, a confidence interval can be placed around point forecasts of out-of-sample outcomes.

Outcomes can be predicted for hypothetical treatments and selection criteria by simply including them in θ_{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.

VII.B Out-of-Population Prediction

An out-of-population prediction uses the estimated model in a different social environment, denoted ν^\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 Λ^\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 ν^\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 λ^\perp . This requires a vector y^\perp of moments drawn from the experimental group with at least $K - 1$ elements. In the worst 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 7. Out-of-population Environment. Given an estimated social environment $\hat{\nu}$, an *out-of-population environment* $\hat{\nu}^\perp \left(Y^\perp, y^\perp, \Psi_p^\perp, \mathcal{H}^\perp [y^\perp; \theta_{\text{cond}}], \hat{\nu} \right)$ is defined as:

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

where

$$\begin{aligned} \hat{\theta}^\perp &\equiv \left(0 \quad 0 \quad 0 \quad 0 \quad 0 \quad \theta_{\text{end}} \quad \hat{\theta}_{\text{exog}}^\perp \quad \Psi_p^\perp \right) \\ \hat{\theta}_{\text{exog}}^\perp &= \left(\hat{\Lambda}^\perp \quad \hat{\Gamma} \right) \\ \hat{\Lambda}^\perp &= \arg \min_{\Lambda^\perp \in \Xi^\kappa} \check{\Delta}(\Lambda^\perp) A^\perp \check{\Delta}(\Lambda^\perp)' \\ \Psi_x^\perp &= \left(0 \quad \mathcal{H}^\perp [y^\perp; \theta_{\text{cond}}] \quad \cdot \quad \cdot \quad \cdot \right). \end{aligned} \tag{63}$$

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 there no exogenous

variation needs to be inserted. (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 policy and demographic information parameters within a model is to interact the in-sample index d directly with the exogenous parameter vector Γ . 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 .

VII.C Randomization and Attrition Bias

Randomization bias occurs when individuals agree to participate in the 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 α contain a choice variable j which equals one when the subject agrees to participate in the next period. Thus, for $t \geq t_0$, $\mathcal{H}[y; \theta_{\text{cond}}] = j$. 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:

$$u^*(\alpha, \theta) = j(u(\alpha, \theta) - \kappa) + (1 - j)u(\alpha, \theta) \Big|_{f=F} \quad (64)$$

where $\kappa \in \Upsilon$ and $u^*(\alpha, \theta)$ 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 κ . 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 δ) are more likely to drop out conditional on the realized state.

To control for randomization bias requires 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 be an entry period as defined above. The extended utility u^* 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 () 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^* .)

VII.D 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 () which contains the Boolean component $\mathcal{B}[y_n(\theta_{\text{cond}}) = Y(\alpha, \theta)]$. 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:

$$\theta_{\text{exog}} \equiv (\Lambda \quad \Gamma \quad \Sigma) \tag{65}$$

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

$$\exp \left\{ -\frac{1}{2} \left(\frac{y_n(\theta_{\text{cond}}) - Y(\alpha, \theta)}{\Sigma} \right) \left(\frac{y_n(\theta_{\text{cond}}) - Y(\alpha, \theta)}{\Sigma} \right)' \right\}. \tag{66}$$

(The division by Σ 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.

VII.E 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: $a' = a^*(\mathcal{B}[a < A] a + 1, [0], [\emptyset])$. The value of terminal states, $V(\theta|_A)$ and the jump values of initial states, $P(\theta|_{a=1})$, 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 single 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, $d' = d^*(d, [\pi_d], [\{d\}])$. 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.

VII.F 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 Ψ_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 \(2002a\)](#) 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. 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 inter-

preted without a model of subject response to the experiment. This preconception has been challenged by Hotz et. al (1999), [Ichimura and Taber\(2000\)](#), and [Rosenzweig and Wolpin \(2000\)](#) on both practical and theoretical grounds. This paper folds their 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 alternative to these unappealing assumptions, a constructive set of tools are 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.

IX. Appendix

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 $z \equiv \arg \max_{\theta_{\text{exp}}} \bar{r}(\theta_{\text{exp}})$ 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 $t = 0$: $t_0(\theta_{\text{exp}}) \equiv -[\bar{r}(\bar{\theta}_{\text{clock}}(z)) - \bar{r}(\theta_{\text{exp}})]$.
- 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 $f' = f^*(1, 0, 0)$ for $(\alpha, \theta) \in h_{\bar{t}}$, and $\bar{t} = \min_e t_{\text{min}}(e) - 1$.

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 (7) for all θ_{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 $t = t_{\text{min}}$. Compute $E_{-\infty}[Y | d]$, $\omega(k; \theta_{\text{cond}})$ and $\Omega\{\theta|k, \theta_{\text{cond}}\}$.
- j.** Increase t by 1. Update Ω and ω . 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 $d = 1$.

X. References

- Aguirregabiria, V. and P. Mira. 2002. "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models," forthcoming, *Econometrica*, July.
- 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 2002. "Finite Mixture Distribution, Sequential Likelihood, and the EM Algorithm," manuscript, Duke University.
- 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, Z. and Wolpin, K. I. 1999. "Why Youths Drop out of High School: The Impact of Preferences, Opportunities, and Abilities," *Econometrica* 67, 6, pp. 1295-1339.
- Ferrall, C. 2002a. "Solving Finite Mixture Models in Parallel," manuscript, Queen's University, <http://qed.econ.queensu.ca/pub/faculty/ferrall/papers/mixtures.pdf>.
- Ferrall, C. 2002b. "Explaining and Forecasting Results of the Self-Sufficiency Project," research report, Social Research Demonstration Corporation, <http://www.srdc.org>.
- 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>.
- 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.
- Meyer, B. D. 1995. "Lessons from the U.S. Unemployment Insurance Experiments," *Journal of Economic Literature* 33, 1, pp. 91-131.
- 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-??.

Shearer, B. 2001. "Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment," working paper, Department of Economics, Laval University, no. 0101, <http://www.ecn.ulaval.ca/w3/recherche/cahiers/2001/0101.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.