IDENTIFICATION AND DECISIONS WITH SOCIAL INTERACTIONS

Charles F. Manski
Department of Economics and Institute for Policy Research, Northwestern University
Leverhulme Visiting Profesor, University College London

Cemmap Masterclass: 17-18 February 2011

This masterclass examines identification of treatment response with social interactions, and studies the implications for decision making. The identification analysis emphasizes partial identification under weak assumptions. The study of decision making considers both social planning and private decisions.

17 February: Identification of Treatment Response

Lecture 1. Shape Restrictions on Response Functions

Lecture 2. Restrictions Derived from Models of Endogenous Interactions

Lecture 3: Distributional Assumptions

18 February: Decision Making with Partial Knowledge of Treatment Response

Lecture 4. Social Planning

Lecture 5: Mechanism Design

Lecture 6: Private Decisions with Social Learning

Lecture Sources

Lectures 1 through 3

Manski, C. (2011), “Identification of Treatment Response with Social Interactions,” Department of Economics, Northwestern University.

Lecture 4

Manski, C. (2010), “Vaccination with Partial Knowledge of External Effectiveness,” *Proceedings of the National Academy of Sciences*, 107, 3953-3960.

Lecture 5

Manski, C. and J. Mayshar (2003), “Private Incentives and Social Interactions: Fertility Puzzles in Israel,” *Journal of the European Economic Association*, 1, 181-211.

Lecture 6

Manski, C. (2004), “Social Learning from Private Experiences: The Dynamics of the Selection Problem,” *Review of Economic Studies*, 71, 443-458.

IDENTIFICATION OF TREATMENT RESPONSE WITH SOCIAL INTERACTIONS

Charles F. Manski
Northwestern University

http://faculty.wcas.northwestern.edu/~cfm754/treatment_with_social_interactions.pdf

Abstract

This paper studies identification of potential outcome distributions when treatment response may have social interactions.

Defining a person's treatment response to be a function of the entire vector of treatments received by the population, I study identification when nonparametric shape restrictions and distributional assumptions are placed on response functions.

An early key result is that the traditional assumption of individualistic treatment response is a polar case within the broad class of *constant treatment response* (CTR) assumptions, the other pole being unrestricted interactions.

Important non-polar cases are interactions within reference groups and anonymous interactions.

I first study identification under assumption CTR alone.

I then strengthen this assumption to semi-monotone response.

I next discuss derivation of these assumptions from models of endogenous interactions.

Finally, I combine assumption CTR with statistical independence of potential outcomes from realized *effective treatments*. The findings both extend and delimit the classical analysis of randomized experiments.

Basic Concepts and Notation

J is a population. (J, Ω, P) is a probability space.

T is a set of feasible treatments. Let $T^J \equiv \times_{j \in J} T$.

For each $j \in J$, response function $y_j(\cdot): T^J \rightarrow Y$ maps treatment vectors into potential outcomes.

Thus, $y_j(t^J)$ is the outcome for j under $t^J \equiv (t_k, k \in J)$.

Person j has realized treatment z_j and outcome $y_j \equiv y_j(z^J)$.

Observation of $[(y_j, z_j), j \in J]$ reveals $P(y, z)$, hence $P[y(z^J)]$.

The objective is to learn $P[y(t^J)]$, $t^J \in T^J$.

Effective Treatments

For each j, let $c_j(\cdot)$: $T^J \to C_j$ be a known function mapping treatment vectors onto a set C_j.

Assumption CTR (Constant Treatment Response):

$$c_j(t^J) = c_j(s^J) \;\Rightarrow\; y_j(t^J) = y_j(s^J).$$

The elements of C_j are the *effective treatments* for person j.

Identification with Assumption CTR

Suppose that one observes $[c_j(\cdot), y_j, z_j; j \in J]$.

One can infer $y_j(t^J)$ if and only if $c_j(z^J) = c_j(t^J)$.

When this occurs, z^J and t^J are effectively the same treatment, yielding the same outcome $y_j(t^J) = y_j(z^J) = y_j$.

When $c_j(z^J) \neq c_j(t^J)$, assumption CTR and observation of y_j do not reveal $y_j(t^J)$.

By the Law of Total Probability,

$$P[y(t^J)] = P[y(t^J)|c(z^J) = c(t^J)] \cdot P[c(z^J) = c(t^J)] + P[y(t^J)|c(z^J) \neq c(t^J)] \cdot P[c(z^J) \neq c(t^J)].$$

Observation reveals $P[c(z^J) = c(t^J)]$ and $P[c(z^J) \neq c(t^J)]$.

CTR gives $P[y(t^J)|c(z^J) = c(t^J)] = P[y|c(z^J) = c(t^J)]$.

Observation reveals $P[y|c(z^J) = c(t^J)]$ if $P[c(z^J) = c(t^J)] > 0$.

Observation and CTR do not reveal $P[y(t^J)|c(z^J) \neq c(t^J)]$.

Proposition CTR: Given Assumption CTR, the identification region for $P[y(t^J)]$ is

$$H\{P[y(t^J)]\} = \{P[y|c(z^J) = c(t^J)] \cdot P[c(z^J) = c(t^J)] + \delta \cdot P[c(z^J) \neq c(t^J)], \delta \in \Delta_Y\}. \qquad \square$$

Interactions within Reference Groups

Let $G(j) \subset J$ denote the reference group for j. An effective treatment is the sub-vector of treatments in the group.

Let $T^{G(j)} \equiv \times_{k \in G(j)} T$, and let $t^{G(j)} \equiv [t_k, k \in G(j)]$.

Let $C_j = T^{G(j)}$ and $c_j(t^J) = t^{G(j)}$.

$$H\{P[y(t^J)]\} = [P(y|z^G = t^G) \cdot P(z^G = t^G) + \delta \cdot P(z^G \neq t^G), \quad \delta \in \Delta_Y].$$

Unrestricted Interactions: G(j) = J for all j.

$$H\{P[y(t^J)]\} = [P(y|z^J = t^J)\cdot P(z^J = t^J) + \delta\cdot P(z^J \neq t^J),\ \delta \in \Delta_Y].$$

Thus, observation is uninformative about the outcome distribution with a counterfactual treatment vector.

Individualistic Response: G(j) = j for all j ∈ J.

$$H\{P[y(t^J)]\} = [P(y|z = t)\cdot P(z = t) + \delta\cdot P(z \neq t),\ \delta \in \Delta_Y].$$

Anonymous Interactions

An interaction is *anonymous* if the outcome of person j is invariant with respect to permutations of the treatments received by other members of his group.

Let G(j)/j denote the reference group exclusive of person j himself and let $\pi[t^{G(j)/j}]$ denote the set of permutations of treatment vector $t^{G(j)/j}$. Then

$$H\{P[y(t^J)]\} = \{[P(y|z = t, z^{G/} \in \pi(t^{G/})]\cdot P[z = t, z^{G/} \in \pi(t^{G/})] + \delta\cdot P(z \neq t \text{ or } z^{G/} \notin \pi(t^{G/}), \delta \in \Delta_Y\}.$$

Distributional Interactions

A *distributional* interaction supposes that response is invariant with respect to the size of the reference group and to permutations of the treatments received by other members of the group.

Let Δ_T denote the space of all distributions on T. For $t^J \in T^J$, let $Q(t^{G(j)/j})$ be the empirical distribution of the treatments in $t^{G(j)/j}$. Then

$$H\{P[y(t^J)]\} =$$

$$\{[P(y|z = t, Q(z^{G/}) = Q(t^{G/})]\cdot P[z = t, Q(z^{G/}) = Q(t^{G/})]$$

$$+ \delta\cdot P(z \neq t \text{ or } Q(z^{G/}) \neq Q(t^{G/}), \delta \in \Delta_Y\}.$$

Semi-Monotone Treatment Response

Let C_j be partially ordered. For $(c, c') \in C_j \times C_j$, either $c < c'$ or $c > c'$ or $c \varnothing c'$.

Let Y be a subset of the real line. Let t^J and s^J be two treatment vectors.

Assumption SMTR:

$$c_j(t^J) \geq c_j(s^J) \;\Rightarrow\; y_j(t^J) \geq y_j(s^J). \qquad \square$$

Let $y_0 \equiv \inf Y$ and $y_1 \equiv \sup Y$.

The empirical evidence and Assumption SMTR yield:

$$c_j(t^J) < c_j(z^J) \Rightarrow y_0 \leq y_j(t^J) \leq y_j$$
$$c_j(t^J) = c_j(z^J) \Rightarrow y_j(t^J) = y_j$$
$$c_j(t^J) > c_j(z^J) \Rightarrow y_j \leq y_j(t^J) \leq y_1$$
$$c_j(t^J) \varnothing c_j(z^J) \Rightarrow y_0 \leq y_j(t^J) \leq y_1.$$

Let $y_{jL}(t^J)$ and $y_{jU}(t^J)$ denote the lower and upper bounds on $y_j(t^J)$.

Proposition SMTR: Given Assumption SMTR, the identification region for $P[y(t^J)]$ is

$$H\{P[y(t^J)]\} = \{\delta \in \Delta_Y: P[y_U(t^J)] \geq_{sd} \delta \geq_{sd} P[y_L(t^J)]\}. \quad \square$$

Let D be any parameter of the outcome distribution that respects stochastic dominance. Then

$$D[y_L(t^J)] \le D[y(t^J)] \le D[y_U(t^J)].$$

For example,

$$y_0 \cdot P[c(t^J) < c(z^J) \cup c(t^J) \varnothing c(z^J)] + E[y \mid c(t^J) \ge c(z^J)] \cdot P[c(t^J) \ge c(z^J)]$$

$$\le E[y(t^J)]$$

$$\le y_1 \cdot P[c(t^J) > c(z^J) \cup c(t^J) \varnothing c(z^J)] + E[y \mid c(t^J) \le c(z^J)] \cdot P[c(t^J) \le c(z^J)].$$

Reinforcing Interactions

Let T be partially ordered.

Let j have reference group G(j) and let $T^{G(j)}$ inherit the partial ordering on T. That is, let $c_j(t^J) \geq c_j(s^J)$ mean that $[t_k \geq s_k, \text{ all } k \in G(j)]$.

A reinforcing interaction occurs when

$$[t_k \geq s_k, \text{ all } k \in G(j)] \Rightarrow y_j(t^J) \geq y_j(s^J).$$

Example: Vaccination of person j against an infectious disease may reduce the chance that he will become ill. Vaccination of others may also reduce his probability of illness, reinforcing the effect of own vaccination.

Monotone Metric Interactions

Let persons be positioned spatially on social networks.

Suppose that the strength of interaction between two persons decreases with the distance between them.

The ordinal essence of this idea may be formalized as a case of semi-monotone treatment response.

Partially order G(j) in terms of distance from j. Let G*(j) denote the partially-ordered version of G(j).

Compare response to certain permutations of a specified treatment vector. Specifically, let two group members who differ in distance from j exchange treatments.

Whereas the person closer to j originally was to receive the smaller of their two treatments, the exchange makes this person receive the larger treatment.

A *monotone metric interaction* occurs if this re-allocation weakly increases the outcome experienced by j.

Let $s^{G*(j)}$ be a specified vector of reference-group treatments.

Let $c_j(t^J) \geq c_j(s^J)$ mean that $t^{G*(j)}$ is a permutation of $s^{G*(j)}$ that exchanges the treatments of two ordered group members, say k and m, with $k < m$, $s_k < s_m$, $t_k = s_m$, and $t_m = s_k$.

A monotone metric interaction occurs if $y_j(t^J) \geq y_j(s^J)$.

Example: Consider a geographic region subject to air pollution created by residential burning of fossil fuels.

Let treatments provide incentives for use of clean fuels.

Let the outcomes of interest be the health status of residents.

One may find it credible to assume that treatment interactions are reinforcing and monotone metric.

Reinforcing means that region-wide strengthening of incentives improves the health status of all residents.

Monotone metric means that the health status of person j improves with a reallocation of heterogeneous incentives within the region, strengthening them for persons who live close to j and correspondingly weakening them for persons who live far from j.

The assumption of a reinforcing interaction is transparent. Strengthening of incentives should induce all residents to use cleaner fuels and, hence, reduce pollution region-wide. If health status decreases with exposure to pollution, the result should be a region-wide increase in health.

The monotone-metric assumption is more subtle. The reallocation of incentives should induce use of cleaner fuels by residents who live close to j and dirtier fuels by those who live far from j. Suppose that pollution decays with distance from the source of the burning. Suppose that persons who live at different distances from j respond similarly to incentives. Then the reallocation should yield a net reduction in the exposure of person j to pollution.

Illustration: Vaccination Against Infectious Disease

Let $T = \{0, 1\}$, with $(\tau = 1)$ denoting vaccination and $(\tau = 0)$ no vaccination.

Let the outcome measure health status, with $y = 1$ if a person is in good health and $y = 0$ if he is ill.

Sufficient statistics for $P(y, z)$ are

$P_{11} \equiv P(y = 1|z = 1)$, $P_{10} \equiv P(y = 1|z = 0)$, $p \equiv P(z = 1)$.

Consider a treatment vector t^J that increases the population rate of vaccination from p to $q > p$. In particular, $t_j = 1$ for all persons with $z_j = 1$ and for some with $z_j = 0$.

The objective is to learn $P[y(t^J) = 1]$.

Individualistic Response

$$H\{P[y(t^J) = 1]\} = [P(y = 1 | z = t)\cdot P(z = t), \quad P(y = 1 | z = t)\cdot P(z = t) + P(z \neq t)].$$

$P(z = t) = p + 1 - q$. $P(z \neq t) = q - p$.

$$P(y = 1 | z = t) = P_{11}[p/(p + 1 - q)] + P(y = 1 | z = t, z = 0)\cdot[(1 - q)/(p + 1 - q)].$$

$P(y = 1 | z = t, z = 0)$ is revealed by the empirical evidence.

Monotone-Individualistic Response

$$H\{P[y(t^J) = 1]\} = [P(y = 1),\ q - p + P(y = 1 \mid t = z)\cdot(p + 1 - q)].$$

The lower bound is larger than using assumption ITR alone. The upper bound is the same as with assumption ITR alone.

Reinforcing Interactions

$$H\{P[y(t^J) = 1]\} = [P(y = 1),\ 1].$$

The lower bound is the same as with monotone-individualistic response. The upper bound is 1 because a reinforcing interaction permits the possibility that increasing the vaccination rate completely eliminates disease transmission.

Derivations of Assumptions CTR and SMTR from Models of Endogenous Social Interactions

I have viewed a response function as a primitive. I posed CTR and SMTR as direct restrictions on this function.

Researchers often model the social mechanism mapping treatments into outcomes.

Economists relate outcomes to choices. They suppose that choices express individual optimizing behavior and the equilibria of games.

Epidemiologists study models of infection and contagion.

From the perspective of such models, response functions are quantities whose properties stem from the mechanism under study. Hence, CTR and SMTR should be derived.

The primitive in a model of endogenous interactions is a system of *structural equations* that takes the outcome of each person to be a function of population treatments and outcomes:

$$y_j(t^J) = f_j[t_j, t^{J/j}, y^{J/j}(t^J)], \quad j \in J.$$

The *structural function* $f_j(\cdot)$ permits $y_j(t^J)$ to be determined by j's own treatment as well as by the treatments and outcomes of other members of the population.

The term *exogenous* interaction describes $t^{J/j}$ as an argument of $f_j(\cdot)$. The term *endogenous* interaction describes $y^{J/j}(t^J)$.

If $y^{J/j}(t^J)$ were not an argument, $f_j(\cdot)$ would be the person's response function. The presence of $y^{J/j}(t^J)$ yields a system of simultaneous equations.

Example: Consider illness from infectious disease. Let the outcome measure health status. Let the treatment be vaccination status. Illness may depend on a person's own vaccination status, on the status of others (exogenous interaction), and on the illness outcomes of others (endogenous interaction).

Example: Consider labor supply in a population of couples. Let the outcome of interest be hours worked. Let the treatment be a person's market wage. One may think it reasonable to assume that labor-supply interactions occur only within couples, not between them. Within each couple, a person's labor supply may depend on his or her own wage, spouse's wage (exogenous interaction), and spouse's labor supply (endogenous interaction).

An outcome vector $y^J(t^J) \equiv [y_j(t^J), j \in J]$ that solves the equations is a *reduced form*.

A model is *complete* if there exists a unique solution for all feasible structural functions. A model is *incomplete* if there are multiple solutions or no solutions.

Identification of Structural and Response Functions

Observation reveals

$$y_j = f_j(z_j, z^{J/j}, y^{J/j}), \qquad j \in J.$$

Econometricians have studied identification of structural function when this evidence is combined with shape restrictions and distributional assumptions on f^J.

We want to use endogenous-interactions models to identify response functions, not structural functions.

A model has identifying power if the empirical evidence and assumptions on f^J imply restrictions on $y^J(t^J)$.

Inference on structural functions has been the dominant theme of econometric research

Econometricians have occasionally observed that the objective may be to infer response functions.

Goldberger put it this way in his 1989nET Interview:

“Well, that's one position, that the entire content in a structural model is simply in the restrictions, if any, that it implies on the reduced form—that’s true. That gives priority to the reduced form.”

The relationship between identification of structural and response functions is straightforward when the structural functions are linear in treatments and outcomes.

Then solution of the structural equations shows that response functions are linear in treatments. The parameters of response functions are many-to-one functions of the parameters of the structural functions. Hence, identification of response functions is a simpler objective than identification of structural functions.

Outside of linear models, the relationship between identification of structural and response functions is largely an open question.

The Linear-in-Means Model

Let the population partition into symmetric reference groups characterized by values for an observed covariate x. Each group contains a continuum of persons.

The linear-in-means model assumes

$$y_j(t^J) = \alpha + \beta_1 t_j + \beta_2 E(t|x_j) + \gamma E[y(t^J)|x_j] + u_j.$$

Taking expectations conditional on x_j yields

$$E[y(t^J)|x_j] = \alpha + (\beta_1 + \beta_2)E(t|x_j) + \gamma E[y(t^J)|x_j] + E(u|x_j).$$

Unless $\gamma = 1$, the unique equilibrium value of $E[y(t^J)|x_j]$ is

$$E[y(t^J)|x_j] = \frac{\alpha}{1-\gamma} + \frac{\beta_1 + \beta_2}{1-\gamma} E(t|x_j) + \frac{E(u|x_j)}{1-\gamma}.$$

Insertion of the right-hand side into the structural function yields the response function

$$y_j(t^J) = \frac{\alpha}{1-\gamma} + \beta_1 t_j + \frac{\gamma\beta_1 + \beta_2}{1-\gamma} E(t|x_j) + \frac{\gamma}{1-\gamma} E(u|x_j) + u_j.$$

The model thus far does not pin down the structural or response functions. The reason is that it does not restrict the unobserved covariates (u_j, $j \in J$).

Assume that $E(u|z, x) = 0$. This implies a linear mean regression relating realized treatments and outcomes:

$$E(y|z, x) = \frac{\alpha}{1 - \gamma} + \beta_1 z + \frac{\gamma\beta_1 + \beta_2}{1 - \gamma} E(z|x) \equiv \varphi_0 + \varphi_1 z + \varphi_2 E(z|x).$$

Observation of realized treatments and outcomes reveals $E(y|z, x)$ on the support of (z, x). Hence, φ is point-identified if the support of $[1, z, E(z|x)]$ is not contained in a linear subspace of R^3.

Knowledge of φ and the empirical evidence imply knowledge of $(u_j, j \in J)$. Knowledge of φ and $(u_j, j \in J)$ implies knowledge of the response functions $[y_j(\cdot), j \in J]$.

Point-identification of the response-function parameters φ does not imply point-identification of the structural parameters $(\alpha, \beta_1, \beta_2, \gamma)$.

β_1 is point-identified but $(\alpha, \beta_2, \gamma)$ are not. Thus, one cannot distinguish exogenous from endogenous interactions under the maintained assumptions.

Nevertheless, the assumptions fully reveal the population vector of response functions.

Complete and Incomplete Models

Given a complete model, identification of $P[y(t^J)]$ is logically no more difficult than identification of f^J.

Let Φ denotes the identification region for f^J. For $f^J \in \Phi$, Let $y^J(t^J, f^J)$ denote the unique solution.

The identification region for $y^J(t^J)$ is $[y^J(t^J, f^J), f^J \in \Phi]$. The cardinality of this set cannot be larger than that of Φ.

Knowledge of $y^J(t^J)$ implies knowledge of $P[y(t^J)]$. Hence, identification of $P[y(t^J)]$ is logically no more difficult than identification of f^J.

Suppose that the model is incomplete, with at least one solution for every feasible value of f^J and multiple solutions for some values.

For each $f^J \in \Phi$, let $Y(t^J, f^J)$ denote the set of solutions.

The identification region for $y^J(t^J)$ is $\{Y(t^J, f^J), f^J \in \Phi\}$.

The cardinality of this set may be larger or smaller than that of Φ. It is larger when the model point-identifies f^J.

Then f^J is known, but $Y(t^J, f^J)$ contains multiple elements. Hence, $H\{P[y(t^J)]\}$ may contain multiple elements.

Structural and Response Reference Groups

Endogenous interactions models assume that interactions occur within known structural reference groups.

A model may assume that

$$y_j(t^J) = f_j[t_j, t^{F(j)/j}, y^{F(j)/j}(t^J)], \quad j \in J.$$

There exists no universal relationship between the structural group F(j) and the response group G(j).

I give three illustrative polar cases. I suppose throughout that the endogenous-interactions model is complete.

Symmetric Structural Groups

Let F be a group of persons and suppose that membership in structural group F is symmetric.

That is, $F(j) = F$ for all $j \in F$.

Then the structural equations pertaining to persons in F are

$$y_j(t^F) = f_j[t_j, t^{F/j}, y^{F/j}(t^F)], \quad j \in F.$$

Completeness implies that these equations have a unique solution $y_j(t^F)$, $j \in F$.

Hence, the members of F share the same response reference group, namely $G(j) = F$ for all $j \in F$.

Recursive Structural Groups

Let J be an ordered set of persons, indexed by the positive integers. Let the structural equations have the form

$$y_1(t^J) = f_1(t_1),$$
$$y_j(t^J) = f_j[t_j, t_{j-1}, y_{j-1}(t^J)], \quad j \in (2, \ldots .).$$

The structural group for person j is F(j) = (j − 1, j).

The response group for j is G(j) = (1, . . . j).

Partly Responsive Structural Groups

Say that a person is *responsive* to treatment if the value of his structural function may vary with his own treatment, all else equal. He is *unresponsive* otherwise. Assume no one is responsive to the treatments of others.

Let R denote the set of responsive persons. Then

$$y_j(t^J) = f_j[t_j, y^{J/j}(t^J)], \quad j \in R,$$

$$y_j(t^J) = f_j[y^{J/j}(t^J)], \quad j \notin R.$$

The structural group is F(j) = J if $j \in R$; F(j) = J/j if $j \notin R$.

The response group is G(j) = R for all $j \in J$.

Observe how the relationship between structural and response reference groups differs across these cases.

Structural and response groups are identical in a population with symmetric structural groups.

Structural groups are smaller than response ones when structural groups are recursive. The structural group of person j is (j − 1, j) and the response group is (1, . . . , j).

Structural groups are larger than response ones when structural groups are partly responsive. The former group is either J or J/j, but the latter is R.

Reinforcing Structural and Response-Function Interactions

A researcher posing an endogenous-interactions model may assume that structural functions are monotone in their arguments.

Lazzati (2010) studies the implications of such assumptions for response functions.

Lazzati supposes that structural groups are symmetric. Hence, structural and response groups are identical. She also supposes that the range space of outcomes is compact.

She considers two forms of monotonicity.

First, suppose that endogenous interactions are reinforcing.

Formally, consider persons in a symmetric structural group F. For all $j \in F$ and $t^F \in T^F$, assume that $f_j(t_j, t^{F/j}, \cdot)$ weakly increases in its final argument, $y^{F/j}(t^F)$.

Tarski's fixed-point theorem to show that the structural equations have at least one solution.

If they have multiple solutions, there exist smallest and largest solutions, whose values may depend on t^F.

Next, suppose that the structural functions are monotone in own treatment and also that exogenous interactions are reinforcing.

Formally, for $j \in F$ and $y^{F/j}(t^F) \in Y^{F/j}$, assume $f_j[\cdot, \cdot, y^{F/j}(t^F)]$ weakly increases in $(t_j, t^{F/j})$.

Combining this monotonicity assumption and the earlier one shows that the smallest and largest solutions to the structural equations are weakly increasing functions of $(t_j, t^{F/j})$.

When the model is complete, these two results imply that response functions satisfy assumption SMTR, with $t^F \geq s^F \Rightarrow y_j(t^F) \geq y_j(s^F)$ for all $j \in F$.

This conclusion need not hold when the model is incomplete. Then the social mechanism at work in the population may possibly select a smaller solution under treatment vector t^F than under s^F.

Lazzati introduces the further assumption that the mechanism always selects either the smallest or the largest solution to the equations.

Then assumption SMTR holds even when the endogenous-interactions model is incomplete.

A caveat is that the credibility of the further assumption may be difficult to assess in applications.

Statistical Independence of Potential Outcomes and Realized Effective Treatments

Assumptions CTR and SMTR restrict the shape of individual response functions, without constraining the distribution of response across the population.

Research under assumption ITR joins shape restrictions on response functions with distributional assumptions.

Similarly, studies of models of endogenous interactions pose shape restrictions and distributional assumptions on structural functions.

A classical union of shape restrictions and distributional assumptions combines assumption ITR with the assumption that potential outcomes are statistically independent of realized treatments.

The statistical independence assumption has high credibility when realized treatments are randomly assigned.

The pair of assumptions transparently yields point identification of potential outcome distributions, provided only that realized treatments equal potential treatments for a positive fraction of the population.

I generalize the classical derivation to settings with social interactions. Specifically, I combine assumption CTR with the assumption that potential outcomes are statistically independent of realized effective treatments.

Analysis with Assumption ITR and SI

Assume that $P[y(\tau)] = P[y(\tau)|z]$.

Assumption ITR implies that $P[y(\tau)|z = \tau] = P(y|z = \tau)$.

Observation reveals $P(y|z = \tau)$ if and only if $P(z = \tau) > 0$.

Hence, $P[y(\tau)]$ is point-identified if and only if $P(z = \tau) > 0$.

Analysis with Assumption CTR and SI

Persons i and j have the same *effective-treatment type* if there exists a permutation operator $\pi_{ij}: T^J \rightarrow T^J$ such that $c_i(t^J) = c_j[\pi(t^J)]$ for all $t^J \in T^J$.

Example: Let i and j have groups of size N. Then $c_i(t^J)$ and $c_j(t^J)$ are subvectors of t^J of length N. A permutation of t^J transforms $c_i(t^J)$ into $c_j(t^J)$.

Let the population be composed of a finite set M of types.

For $m \in M$, let C_m be the common finite set of effective treatments for persons in J_m. Let $J_{m\gamma} \equiv [j \in J_m: c_j(t^J) = \gamma]$.

Assumption SI: For each group $J_{m\gamma}$ with $P(J_{m\gamma}) > 0$,

$$P[y(t^J)|J_{m\gamma}] = P[y(t^J)|J_{m\gamma}, c(z^J)].$$

Assumption CTR implies that $P[y(t^J)|J_{m\gamma}, c(z^J) = \gamma] = P[y|J_{m\gamma}, c(z^J) = \gamma]$.

Observation of realized treatments and outcomes reveals $P[y|J_{m\gamma}, c(z^J) = \gamma]$ if and only if $P[c(z^J) = \gamma|J_{m\gamma}] > 0$.

Hence, Assumption SI point-identifies $P[y(t^J)|J_{m\gamma}]$ if and only if $P[c(z^J) = \gamma|J_{m\gamma}] > 0$.

The Law of Total Probability gives

$$P[y(t^J)] = \sum_{(m \in M,\ \gamma \in C_m)} P[y(t^J) | J_{m\gamma}] \cdot P(J_{m\gamma}).$$

Hence,

Proposition SI: Given assumption SI, the identification region for $P[y(t^J)]$ is

$$H\{P[y(t^J)]\} = \{ \sum_{(m \in M,\ \gamma \in C_m:\ P[c(z^J) = \gamma | J_{m\gamma}] > 0)} P[y | J_{m\gamma}, c(z^J) = \gamma] \cdot P(J_{m\gamma})$$

$$+ \quad \delta \cdot \sum_{(m \in M,\ \gamma \in C_m:\ P[c(z^J) = \gamma | J_{m\gamma}] = 0)} P(J_{m\gamma}), \qquad \delta \in \Delta_Y\}. \quad \square$$

Identifying Power

Assumption SI reveals $P[y(t^J)]$ if and only if $P[c(z^J) = \gamma | J_{m\gamma}] > 0$ for all $m \in M$ and $\gamma \in C_m$ such that $P(J_{m\gamma}) > 0$.

Under assumption ITR, this reduces to $P(z = \tau | J_\tau) > 0$ for all $\tau \in T$ such that $P(J_\tau) > 0$.

This support condition generically holds if realized treatments are assigned randomly with ex ante assignment probabilities $\varphi(\tau) > 0$, all $\tau \in T$.

Then familiar arguments using laws of large numbers show that $P(z = \tau | J_\tau) \cong \varphi(\tau)$ for all τ such that $P(J_\tau) > 0$.

The support condition is more subtle with interactions.

Let persons of type m have reference groups of size S.

Then $C_m = T^S$ and $P[c(z^J) = \gamma | J_{m\gamma}] = P[z^{G(\cdot)} = \gamma | J_{m\gamma}]$.

Thus, equalizing the realized and conjectured effective treatments of person j necessitates fixing the realized treatments of all members of his reference group G(j).

For each $k \in G(j)$, let $D(k) \equiv [i \in J: k \in G(i)]$ denote the subset of J who list k as a member of their groups.

Equalizing the realized and conjectured effective treatments of person j constrains the realized effective treatments of the entire class of persons $[D(k), k \in G(j)]$.

This can make it difficult to satisfy the support condition when (a) reference groups are large or (b) there exist persons who belong to many reference groups.

I give two illustrations. In both cases, interactions are *global* in the sense that the class of persons $[D(k), k \in G(j)]$ comprises the entirety of J_m.

Groups with Leaders and Followers

Let type m be persons with reference groups of size N + L.

Membership is symmetric for N persons.

Membership is asymmetric for L persons. These persons are in all groups but are not themselves type-m.

Call the L persons *leaders* and the N persons *followers*.

The effective treatment of a person of type m is a subvector of t^J of length N + L.

Let $\Lambda(m) \subset J$ denote the leaders of type m.

Let $\gamma = (\tau^N, \tau^{\Lambda(m)})$ denote a situation in which followers receive the N treatments τ^N and leaders receive the L treatments $\tau^{\Lambda(m)}$. Then

$$\begin{aligned} P[c(z^J) = \gamma | J_{m\gamma}] &= P[(z^N, z^{\Lambda(m)}) = (\tau^N, \tau^{\Lambda(m)}) | J_{m\gamma}] \\ &= P(z^N = \tau^N | J_{m\gamma}) \cdot 1[z^{\Lambda(m)} = \tau^{\Lambda(m)}]. \end{aligned}$$

Thus, $P[c(z^J) = \gamma | J_{m\gamma}] > 0$ if and only if

(a) $z^{\Lambda(m)} = \tau^{\Lambda(m)}$

(b) $P(z^N = \tau^N | J_{m\gamma}) > 0$.

Population-wide Distributional Interactions

Let there exist one type. The reference group is the entire population. Interactions are distributional.

Then $c_j(t^J) = [t_j, Q(t^J)]$ and $c_j(z^J) = [z_j, Q(z^J)]$ for all $j \in J$.

The feasible values of γ are the pairs $[\tau, Q(t^J)]$, $\tau \in T$.

Fixing τ and letting J_τ be the subpopulation who would receive τ under potential treatment vector t^J,

$$\begin{aligned} P[c(z^J) = \gamma | J_{m\gamma}] &= P\{[z, Q(z^J)] = [\tau, Q(t^J)] | J_\tau\} \\ &= P(z = \tau | J_\tau)\cdot 1[Q(z^J) = Q(t^J)]. \end{aligned}$$

Thus, $P[c(z^J) = \gamma | J_{m\gamma}] > 0$ if and only if

(a) $Q(z^J) = Q(t^J)$

(b) $P(z = \tau | J_\tau) > 0$.

Random Assignment of Realized Treatments

Random assignment may motivate assumption CTR-SI. However, it may not have identifying power.

Consider a random assignment process that independently assigns persons to treatments, with ex ante probability distribution π on T. Let π be non-degenerate.

Random assignment does not yield a determinate vector of realized treatments. It yields an ex ante probability distribution for z^J.

The two illustrations demonstrate the difficulty.

Groups with leaders and followers: Observation and Assumption SI-ET can be informative about $P[y(t^J)|J_{m\gamma}]$ only if $z^{\Lambda(m)} = \tau^{\Lambda(m)}$. The ex ante probability that $z^{\Lambda(m)} = t^{\Lambda(m)}$ is $\prod_{j \in \Lambda(m)} \pi(t_j)$. Hence, random assignment yields positive ex ante probability that $z^{\Lambda(m)} \neq t^{\Lambda(m)}$.

Population-wide distributional interaction: Observation and assumption SI-ET can be informative about $P[y(t^J)|J_{m\gamma}]$ only if $Q(z^J) = Q(t^J)$. This equality occurs with ex ante probability less than one. In an uncountably large population, $Q(z^J) = \pi$.

These negative findings do not appear in analysis of random assignment under Assumption ITR, nor in previous efforts to study random assignment in settings with social interactions.

Authors have assumed that the population partitions into a large number of reference groups, each of finite size. Assumption ITR holds when the population is defined to be a collection of groups rather than persons.

The present scenarios have a different structure. They exhibit *global* social interactions rather than *local* ones. When interactions are global, random assignment does not retain its classical identifying power.

Estimation with Sample Data

Draw a random sample of N persons, say J_N, and observe $[c_j(z^J), c_j(t^J), y_j; j \in J_N]$.

One may estimate the identification regions under Assumptions CTR, SMTR, and STR-SI by sample analogs

$$H\{P_N[y(t^J)]\} \equiv \{P_N[y|c(z^J) = c(t^J)]\cdot P_N[c(z^J) = c(t^J)] + \delta\cdot P_N[c(z^J) \neq c(t^J)], \delta \in \Delta_Y\},$$

$$H\{P_N[y(t^J)]\} \equiv \{\delta \in \Delta_Y: P_N[y_U(t^J)] \geq_{sd} \delta \geq_{sd} P_N[y_L(t^J)]\}.$$

$$H\{P_N[y(t^J)]\} = \{ \sum_{(m \in M, \gamma \in C_m: P_N[c(z^J) = \gamma|J_{m\gamma}] > 0)} P_N[y|J_{m\gamma}, c(z^J) = \gamma]\cdot P_N(J_{m\gamma}) + \delta\cdot \sum_{(m \in M, \gamma \in C_m: P_N[c(z^J) = \gamma|J_{m\gamma}] = 0)} P_N(J_{m\gamma}), \quad \delta \in \Delta_Y\}.$$

NOTE: This requires observation of realized effective treatments $[c_j(z^J), j \in J_N]$, not just realized own treatments $(z_j, j \in J_N)$.

Conclusion

This analysis provides the foundation for further work.

Many response-function assumptions beyond CTR, SMTR, and CTR-SI warrant attention.

Further derivation of response-function assumptions from models of social mechanisms would be welcome.

SOCIAL PLANNING UNDER AMBIGUITY

When studying collective decision problems, economists have asked how an optimizing planner should behave.

A standard exercise specifies a set of feasible policies and a social welfare function.

The planner is presumed to know the welfare achieved by each policy.

The objective is to characterize the optimal policy.

In practice, we typically have only partial knowledge of the welfare achieved by alternative policies.

Hence, we cannot determine optimal policies.

This limits the relevance of the standard exercise to actual policy analysis.

A fundamental source of partial knowledge is the identification problem arising from unobervability of counterfactual policy outcomes.

At most one can observe the outcomes that have occurred under realized policies. The outcomes of unrealized policies are logically unobservable.

Yet, determination of an optimal policy requires prospective comparison of all feasible policies.

When economists have studied planning with partial knowledge, it has been standard to assert a subjective distribution over unknown quantities and propose choice of an action that maximizes subjective expected welfare.

This is reasonable when a planner has a credible basis for asserting a subjective distribution on unknown quantities.

However, a subjective distribution is a form of knowledge. A planner may not have a credible basis for asserting one.

Then the planner faces a problem of decision under *ambiguity*.

I have studied how identification problems generate ambiguity about optimal policies.

The planning problems studied share a simple structure.

The task is to choose treatments for a population whose members may vary in their response to treatment.

The welfare function sums the outcomes of the population.

Ambiguity arises when a planner has partial knowledge of treatment response and cannot determine the optimal policy.

Vaccination with Partial Knowledge of External Effectiveness

The problem of choosing an optimal vaccination policy for a population susceptible to infectious disease has drawn considerable attention.

Research studying optimal vaccination has assumed the planner knows how vaccination affects illness rates.

There are two reasons why a planner may have only partial knowledge of the effect of vaccination on illness.

He may not know the *internal* effectiveness of vaccination in generating an immune response that prevents a vaccinated person from become ill or infectious.

He may not know the *external* effectiveness of vaccination in preventing transmission of disease to members of the population who are unvaccinated or unsuccessfully vaccinated.

Knowledge of external effectiveness is most problematic.

A standard randomized clinical trial enables evaluation of internal effectiveness.

However, the outcome data only reveal the external effectiveness of the chosen vaccination rate.

Outcomes with other vaccination rates are counterfactual.

To cope with the absence of empirical evidence, researchers have used epidemiological models to forecast the outcomes that would occur with counterfactual vaccination rates.

Authors typically do not know the accuracy of their assumptions about individual behavior, social interactions, and disease transmission.

Hence, it is prudent to view their forecasts more as computational experiments than as accurate predictions of policy impacts.

Manski (*PNAS*, 2010) studies choice of vaccination policy when a planner has partial knowledge of the external effectiveness of vaccination.

I take the planner’s objective to be minimization of the social cost of illness and vaccination.

I consider a simple scenario in which the population is composed of observationally identical persons.

The planner knows that vaccination is fully effective internally. Thus, vaccinated persons never become ill.

Regarding external effectiveness, the planner knows only that the rate of illness among unvaccinated persons decreases as the vaccination rate rises.

The planner observes the illness rate under a policy that vaccinates an observed fraction of the population.

Dominated Vaccination Rates

I show that the empirical evidence and the assumption of monotone external effectiveness imply that certain vaccination rates are strictly dominated.

That is, there exist other vaccination rates that yield lower social cost whatever the true external effectiveness of vaccination may be.

Broadly speaking, low (high) vaccination rates are dominated when the cost of vaccination is low (high)

Minimax and Minimax-Regret Vaccination Rates

How might the planner choose among the undominated vaccination rates?

I derive the minimax and minimax-regret vaccination rates.

These criteria protect against poor outcomes, but in different ways.

The former chooses an action that minimizes maximum cost across all feasible states.

The latter chooses an action that minimizes maximum regret. The regret of a specified vaccination rate in a given state of nature is the cost of this rate minus the cost of the best possible rate.

I refer to minimax and minimax-regret as “reasonable” decision criteria, not as “optimal” ones.

There is no uniquely correct way to choose among undominated actions.

The crux of the problem in choice under ambiguity is that the planner does not know which action is best.

Optimal Vaccination With Linear External Response

Suppose that the planner must choose the vaccination rate for a large population of observationally identical persons.

Assume that vaccination always prevents a vaccinated person from becoming ill.

Let p(t) be the *external-response function*, giving the fraction of unvaccinated persons who become ill when the vaccination rate is t.

The fraction of the population who become ill is p(t)(1 − t).

The planner wants to minimize a social cost function with two components, the harm caused by illness and the cost of vaccination.

Let $a = 1$ denote the mean social harm caused by illness and let $c > 0$ denote the mean social cost per vaccination, measured in commensurate units. The social cost of vaccination rate t is

$$K(t) = p(t)(1 - t) + ct.$$

The planner wants to solve the problem $\min_{t \in [0, 1]} K(t)$.

Let $p(t) = \rho(1 - t)$ and $0 < \rho \leq 1$.

The optimal vaccination rate is

$$t^* = \underset{t \in [0,1]}{\text{argmin}}\ \rho(1 - t)^2 + ct.$$

The optimal rate is

$$\begin{aligned} t^* &= 0 && \text{if } 2\rho < c. \\ &= 1 - c/(2\rho) && \text{if } 2\rho \geq c. \end{aligned}$$

Partial Knowledge of External Effectiveness

The planner observes the vaccination and illness rates of a study population, whose vaccination rate has been chosen to be some value less than one.

He assumes that the study population and the treatment population have the same external-response function.

He assumes that the illness rate of unvaccinated persons weakly decreases as the vaccination rate increases.

He makes no assumption about the magnitude of the external effect of vaccination.

Let $r < 1$ denote the observed vaccination rate and $q(1 - r)$ denote the observed realized illness rate. The two maintained assumptions are

Assumption 1 (Study Population): The planner observes r and $q(1 - r)$. He knows that $q = p(r)$.

Assumption 2(Vaccination Weakly Prevents Illness): The planner knows that $p(t)$ is weakly decreasing in t.

These assumptions imply that

$$t \le r \;\Rightarrow\; p(t) \ge q,$$
$$t \ge r \;\Rightarrow\; p(t) \le q.$$

Dominance

Let Γ be the set of feasible external-response functions under Assumptions 1 and 2. Let γ index the elements of Γ, the symbol p being reserved for the unknown true function. For $\gamma \in \Gamma$, let

$$S(t, \gamma) \equiv \gamma(t)(1 - t) + ct$$

be the social cost of vaccination rate t when the external-response function is γ.

Rule t is *strictly dominated* if and only if there exists another vaccination rate t' such that $S(t, \gamma) > S(t', \gamma)$ for all $\gamma \in \Gamma$.

A candidate vaccination rate t is strictly dominated if any of these conditions hold:

(a) Let $c < q$. Then t is strictly dominated if $t < r$.

(b) Let $c > q$. Then t is strictly dominated if

$$t > r + q(1 - r)/c.$$

(c) Let $c > 1$. Then t is strictly dominated if

$$(1 - q)/(c - q) < t \leq r \text{ or if } t > \max(r, 1/c).$$

Elimination of dominated rates takes the planner part way toward solution of the vaccination problem.

Decision theory does not provide a consensus prescription for a complete solution.

Instead, it offers alternative criteria that generically ensure choice of an undominated alternative.

Vaccination to Minimize Expected Social Cost

The expected utility criterion recommends that the planner place a subjective distribution on the feasible $p(\cdot)$ and minimize subjective expected social cost. The cost function is linear in $p(\cdot)$, so subjective expected social cost is

$$E_{\Psi}[S(t)] = \pi(t)(1 - t) + ct,$$

where Ψ is the subjective distribution and $\pi(\cdot) \equiv E_{\Psi}[p(\cdot)]$ is the subjective mean of $p(\cdot)$.

The planner acts as a pseudo-optimizer, using the expected external-response function π as if it were the actual external-response function p.

Applications of the expected utility model have typically assumed that the support of a subjective distribution is a finite-dimensional real space.

However, Assumptions 1 and 2 only require that $p(\cdot)$ be an element of an abstract space of weakly decreasing functions.

A planner who assumes that $p(\cdot)$ belongs, with subjective probability one, to a finite-dimensional subset of this function space asserts considerable knowledge beyond Assumptions 1 and 2.

Permitting the support of the subjective distribution to be a larger subset of the function space is mathematically delicate.

Minimax Vaccination

For each candidate vaccination rate t, compute the maximum social cost that can occur across all feasible external-response functions; that is, $\max_{\gamma \in \Gamma} S(t, \gamma)$.

The minimax criterion selects the vaccination rate that minimizes this maximum social cost. Thus, it solves

$$\min_{t \in [0, 1]} \max_{\gamma \in \Gamma} \gamma(t)(1 - t) + ct.$$

The minimax vaccination rate is

$$
\begin{aligned}
t^m &= 0 && \text{if} \quad c > 1 \text{ and } 1 \leq q(1 - r) + cr, \\
&= r && \text{if} \quad c > 1 \text{ and } 1 \geq q(1 - r) + cr \\
& && \text{or if} \quad q < c < 1, \\
&= \text{all } t \in [0, 1] && \text{if } c = q \text{ and } q = 1, \\
&= \text{all } t \in [r, 1] && \text{if } c = q \text{ and } q < 1, \\
&= 1 && \text{if} \quad c < q.
\end{aligned}
$$

Minimax-Regret Rate

For each feasible external-response function $\gamma \in \Gamma$, let

$$S^*(\gamma) \equiv \min_{t \in [0, 1]} S(t, \gamma)$$

denote the smallest social cost achievable when the external-response function is γ. The regret of vaccination rate t in state of nature γ is $S(t, \gamma) - S^*(\gamma)$.

Regret measures the difference between the cost delivered by rate t and that delivered by the best possible rate.

The minimax-regret criterion selects the rate that minimizes maximum regret across all states of nature. Thus, it solves

$$\min_{t \in [0, 1]} \ \max_{\gamma \in \Gamma} S(t, \gamma) - S^*(\gamma).$$

(a) Let $c \leq q$. Then the minimax-regret vaccination rate is

$$t^{mr} = (q + cr)/(q + c).$$

(b) Let $c > q$. Then the minimax-regret vaccination rate is

$$t^{mr} = \underset{t \in [0, 1]}{\text{argmin}}\ 1[t < r]\cdot\{\max [(1 - q)(1 - t), (1 - t) + c(t - r), (c - q)t]\} + 1[t \geq r]\cdot\{\max [q(1 - t), c(t - r), (c - q)t]\}.$$

Numerical Examples

1. Let $c = 0.02$. The planner observes a study population with no vaccination ($r = 0$) and with illness rate $q = 1/5$.

A planner who believes the external-response function is linear would conclude that $\rho = 1/5$ and would choose the vaccination rate $t^* = 7/8$.

A planner who only knows the function to be weakly decreasing would not be able to conclude that any vaccination rates are dominated, because $c < q$ and $r = 0$.

The minimax vaccination rate is $t^m = 1$.

The minimax-regret rate is $t^{mr} = 4/5$.

2. Let $r = 1/2$ and $q = 1/10$. Continue to let $c = 0.02$.

A planner who believes the external-response function is linear would still conclude that $\rho = 1/5$ and choose $t^* = 7/8$.

A planner who only knows the function to be weakly decreasing can determine that any vaccination rate smaller than 1/2 is strictly dominated.

The minimax vaccination rate remains $t^m = 1$.

The minimax-regret rate is $t^{mr} = 5/6$.

3. Let $c = 0.25$. Continue to let $r = 1/2$ and $q = 1/10$.

A planner who believes the external-response function is linear would choose $t^* = 3/8$.

A planner who only knows the function to be weakly decreasing can conclude that any vaccination rate larger than 7/10 is strictly dominated.

The minimax vaccination rate is $t^m = ½$.

The minimax-regret rate is $t^{mr} = ½$.

Related Planning Problems

The analysis extends to settings where vaccination has imperfect but known internal effectiveness.

Population members may have observable covariates.

A planner who cannot mandate vaccination may provide incentives for private vaccination.

In dynamic planning problems, a planner vaccinates a sequence of cohorts, using observation of past outcomes to inform present decisions.

Looking beyond vaccination, the analysis demonstrates how one may address a class of choice problems where a planner observes the outcome of a status-quo policy and feels able to partially extrapolate to counterfactual policies.

Manski (2006) studied the criminal-justice problem of choosing a rate of search for evidence of crime, when a planner has partial knowledge of the deterrent effect of search on the rate of crime commission. I considered a planner who wants to minimize the social cost of crime, search, and punishment. The planner observes the crime rate under a status-quo search rate and assumes that the crime rate falls as the search rate rises.

The formal structure of this planning problem is similar to that of the vaccination problem, the substantive difference between the two notwithstanding.

Mechanism Design under Ambiguity

Manski (*PNAS*, 2010) assumes that the planner has the power to mandate vaccination of the population.

This is realistic in some settings: vaccination of health care workers, military personnel, or students in public schools.

More commonly, vaccination is a private decision, which a public health agency may seek to influence through provision of incentives.

Economists often refer to choice of an incentive policy as a *mechanism design* problem.

Suppose that a planner selects an incentive policy from a set of feasible policies.

Given a policy, members of the population individually choose whether or not to be vaccinated.

Then the resulting social cost may depend not only on the fraction of the population who choose to be vaccinated but also on the composition of the vaccinated group.

The reason is that members of the population may vary in their susceptibility to illness and in the extent to which they can infect other members of the population.

Hence, the effectiveness of an incentive policy may depend on which as well as how many persons choose to be vaccinated.

Public health agencies have limited knowledge of how incentive policies affect private vaccination choices and of the resulting implications for disease transmission.

Hence, a planner choosing an incentive policy faces more ambiguity than does one who mandates vaccination.

Design of incentive policies for vaccination under ambiguity is an open question.

More generally, design of incentive policies under ambiguity is an open question.

I will use child allowance policy to illustrate.

Private Incentives and Social Interactions:

Fertility Puzzles in Israel

Charles F. Manski and Joram Mayshar

(*Journal of the European Economic Association*, 2003)

This paper explores how private and social incentives for fertility may have combined to produce the complex fertility pattern observed in Israel in the past half-century.

Fertility has declined within some ethnic-religious groups, moderately increased in others, and parts of the ultra-orthodox Jewish population have experienced a *reverse fertility transition*, in which childbearing has increased rapidly and substantially.

We present a theoretical analysis of the social dynamics of fertility that shows how private preferences, preferences for conformity to social norms in childbearing, and piecewise linear child allowances could have combined to yield such a complex fertility pattern.

We then explain the identification problem that makes it so difficult to infer the actual Israeli fertility process from data on completed fertility.

Introduction

Discussions of fertility sometimes give the impression that economics and sociology offer competing models of fertility, but the two are compatible.

Women may choose family size to maximize utility functions that recognize both private and social incentives for fertility.

Social incentives may themselves evolve as an outcome of childbearing decisions.

Table 1 summarizes data on *completed fertility* among women alive and married in 1995, whose year of first marriage was either before 1955 or in the period 1970–80.

Table 1. Average Completed Fertility among Women Alive in 1995

	Year of First Marriage			
	Prior to 1955		**1970–1980**	
Ethnicity	Ultra-Orthodox	Others	Ultra-Orthodox	Others
Jew – Mizrahi	6.86 (217)	5.23 (10491)	5.11 (274)	3.46 (20427)
Jew – Ashkenazi	2.76 (571)	2.29 (14954)	5.88 (275)	2.88 (17164)
Jew – Israeli parents	3.03 (40)	3.06 (793)	5.91 (57)	3.12 (2636)
Arab – Non-Bedouin	8.41 (2290)		5.55 (6055)	
Arab – Bedouin	7.32 (34)		9.02 (139)	

In parentheses: the number of women in each cell in the sample.

In a *reverse fertility transition*, fertility increased rapidly and substantially among ultra-orthodox women of non-Mizrahi origins.

What explains the dissimilar levels and trends in completed fertility among different groups of Israeli women?

The potentially relevant sources of time-series and cross-sectional variation in fertility include

The fertility transition

Standard explanations of the fertility transition can rationalize the substantial declines in fertility among Israeli Arabs and among Jews of Mizrahi origin.

In both groups, women whose first marriage occurred prior to 1955 largely lived in traditional societies characterized by low income and high child mortality, but women who first married in 1970-1980 lived in the western-oriented society of modern Israel, characterized by relatively high income and low child mortality.

Religiosity

Ultra orthodox Jewish women follow the directives of their rabbis, who encourage high fertility and discourage the use of contraceptives.

Other Jewish women may hold varying preferences for family size and are not subject to strong religious strictures against contraception.

Social Interactions

The various ethnic and religious groups within Israeli society interacted to different degrees over the past half-century, and so may have been subject to different social norms for childbearing.

Adoption of the fertility norms of Ashkenazi Jews by the Mizrahi Jews, most of who migrated to Israel in the 1950s, may partly explain the fertility decline in the latter group.

Increased integration between the two groups may partly explain the increased fertility among the former.

Yet, throughout the past half-century, the Jewish ultra-orthodox, the Jewish non-ultra-orthodox, and the Arab sub-populations of Israel have largely resided, been schooled, and worked in separate, almost isolated communities.

Recent history

The traumatic decimation of European Jewry in the Holocaust may have affected fertility in the period following World War II.

It is sometimes argued that the Holocaust reduced the fecundity of the women who were directly affected, but increased the desire to have children among survivors and the kin of non-survivors.

It has also been suggested that the relatively high death rates experienced or expected by some sub-groups of Israelis in repeated wars may have increased the desired number of children, as a form of insurance.

Child allowances and related public welfare programs

Changes in public policy have generated time-series variation in private incentives for fertility.

In 1970–75, the Israeli government introduced a universal allowance program.

Modest monthly payments are made to families with one or two children under age 18. The payments for each child from the fourth up are substantial.

Women whose first marriage occurred prior to 1955 bore their children before the introduction of the child allowance programs, while women who first married after 1970 were subject to them throughout their childbearing years.

The puzzle, or set of puzzles, is how to disentangle these forces and identify the socioeconomic process at work.

From a policy perspective, there is particular interest in learning how the child allowance have affected, and continue to affect, fertility.

From a social science perspective, there is a general interest in learning how private and social incentives interact to determine childbearing.

Plan

* Description of Child Allowances and Fertility Patterns

* Model of Family-Size Decisions

* Attempt at Structural Empirical Inference

The Israeli Child Allowance Program

In the tax reform of 1975, Israel quadrupled the child allowances granted up to 1969. Since 1975, the National Insurance Institute (NII) has annually allocated more than 1.6 percent of GDP to this program.

The size of the allowance varies with the number of children, with the first two children under age 18 receiving minimal benefits, and each child from the fourth on receiving a large benefit.

The marginal allowance granted for children from the fourth up fully covers the marginal cost of caring for these children in a manner that averts poverty, as calculated using Israel's official poverty line for families of different sizes.

Consider families with six children under age 18. In 1999, such families received a monthly tax-exempt allowance of NIS 2,566, equivalent to about $640 per month.

Family-Size Decisions: Private Preferences, Child Allowances, and Social Interactions

It is plausible that women choose family size to maximize utility functions that recognize both private and social incentives for fertility.

The private incentives depend in part on government policy toward children, one aspect of which is the child allowance program.

The social incentives depend in part on the childbearing behavior of peers.

Maintained Assumptions

A woman chooses family size irrevocably sometime after marriage, based on the information available at the time.

Let t denote the date of the fertility decision, j denote a woman choosing family size at that date, and k denote family size.

We assume that woman j assigns utility $U_{tj}(k)$ to having k children and maximizes utility. Thus the woman chooses

$$k_{tj} = \text{argmax}_{k = 0, 1, \ldots} U_{tj}(k).$$

We assume no infertility and no child mortality, or similar sources of uncertainty that could prevent some women from realizing their family-size plans. We also take the spacing of the k births as predetermined.

The utility function has three components, two expressing private incentives and the third expressing social incentives:

(i) a private utility of family size, assumed quadratic in the number of children. This is $a_j k - b_j k^2$.

(ii) an incentive derived from the child allowances formula at the time of the fertility decision. This is $A_{tj}(k)$, the discounted life-cycle allowance that woman j would receive under the formula prevailing at the decision date, if j were to have k children.

Let I_{tj} denote family expected life-cycle income from date t on (assume that income does not depend on family size). Then expected life-cycle consumption if woman j were to have k children is $A_{tj}(k) + I_{tj}$. The contribution to utility is $c_j[A_{tj}(k) + I_{tj}]$.

(iii) a social interaction where the woman compares k with the family sizes of others in the population.

The population is composed of M symmetric reference groups.

Consider group m. Woman j incurs a utility loss that grows with the squared deviation between k and the size of each family in group m. The woman averages these losses across the members of m, yielding $\sum_h (k - h)^2 P_{tm}(h)$, where $P_{tm}(h)$ is the fraction of women in group m having h children.

The woman aggregates the average losses across the M groups, yielding the aggregate loss

$$-\sum_m w_{mj} [\sum_h (k - h)^2 P_{tm}(h)],$$

where $w_{mj} \geq 0$, m = 1,…,M.

Combining the private, government, and social determinants, the form of the utility function is

$$U_{tj}(k) = a_j k - b_j k^2 + c_j[A_{tj}(k) + I_{tj}]$$
$$- \textstyle\sum_m w_{mj} [\sum_h (k - h)^2 P_{tm}(h)].$$

Optimal Family Size

Let $S_{tm} = \sum_h h\, P_{tm}(h)$ be the mean number of children in group m at date t. For each m,

$$\sum_h (k - h)^2 P_{tm}(h) = k^2 - 2kS_{tm} + S_{tm}^{\ 2} + \sum_h (h - S_{tm})^2 P_{tm}(h).$$

It follows that

$$U_{tj}(k) = U_{tj}(0) + (a_j + 2\sum_m w_{mj} S_{tm})k$$
$$- (b_j + \sum_m w_{mj})k^2 + c_j A_{tj}(k).$$

Hence

$$k_{tj} = \text{argmax}_{\ k=0,1,\ldots}(a_j + 2\sum_m w_{mj} S_{tm})k$$
$$- (b_j + \sum_m w_{mj})k^2 + c_j A_{tj}(k).$$

The model yields a simple solution if the child allowance is proportional to the number of children in a family. Let $A_{tj}(k) = \pi_{tj}k$ for $\pi_{tj} \geq 0$. Then

$$k_{tj} = \operatorname{argmax}_{k=0,1,\ldots} (a_j + 2\sum_m w_{mj} S_{tm} + c_j\pi_{tj})k - (b_j + \sum_m w_{mj})k^2.$$

Both composite coefficients are positive. Hence utility is quadratic in family size with its maximum at

$$k_{tj} = \text{INT}\left[\frac{a_j + 2\sum_m w_{mj} S_{tm} + c_j\pi_{tj}}{2(b_j + \sum_m w_{mj})} + \tfrac{1}{2}\right].$$

The Social Dynamics of Fertility in Two Special Cases

The fertility model is still too complex to permit much in the way of theoretical analysis.

An interesting analysis becomes possible if we make four further assumptions:

(1) The number of children is continuous.

(2) Each woman is influenced only by women in her own group who made their choices in the preceding period.

(3) All women in a given ethnic-religious group have the same values of utility parameters (b, c, w), so that only the parameter a_j may vary within the group. The within-group distribution of this parameter is continuous and time-invariant, with density f(a).

(4) The child allowance formula is piecewise linear in the number of children, with one kink. The allowance formula is time-invariant and is the same for all members of a given group.

Allowances Proportional to Family Size

Let the allowance formula be $A(k) = \pi k$. Then

$$k_{tj} = \frac{a_j + 2wS_{t-1} + c\pi}{2(b + w)},$$

where S_{t-1} is the mean family size of women who chose their family size at date t – 1. The mean family size of women who choose at date t is

$$S_t = S(S_{t-1}, \lambda) = E\left[\frac{a + 2wS_{t-1} + c\pi}{2(b + w)}\right] = \frac{\quad + 2wS_{t-1} + c\pi}{2(b + w)},$$

where $\mu \equiv E(a) > 0$ and where $\lambda \equiv [b, c, w, A(\cdot)\ f(\cdot)]$ is the full set of parameters that characterize preferences and child allowances.

Steady-State Mean Family Size

With $b > 0$ and $w \geq 0$, S_t converges monotonically to the unique steady state (social equilibrium) solving

$$S^o = S(S^o, \lambda).$$

The unique steady-state is

$$S^o = \frac{\mu + c\pi}{2b}.$$

S^o does not vary with w; hence there is no *social multiplier* effect. The fertility gap between any two groups who differ only in μ is independent of the magnitude of the allowances.

Steady-State Distribution of Family Size

With S^o determined, we can characterize the steady-state distribution of family sizes. The steady-state optimal family size of a woman with private utility parameter a is

$$k^o(a, \lambda) = \frac{a + 2wS^o + c\pi}{2(b + w)} = \frac{1}{2b}\left(\frac{ba + w\mu}{b + w} + c\pi\right).$$

$k^o(a, \lambda)$ is linear in a. Hence the density function for optimal family size has the same shape as f(a), the density of a.

The child-allowance π shifts $k^o(a, \lambda)$ by the same amount, regardless of the value of a. Hence child allowances affect the central tendency of the distribution of family sizes but not its dispersion.

$k^o(a, \lambda)$ is determined by a weighted average of a woman's own utility parameter and the mean μ of this parameter in the group.

Thus, w affects the dispersion of the distribution. All else equal, the larger the value of w, the more concentrated is the distribution around S^o.

Piecewise Linear Child Allowances

Let the allowance formula have the piecewise linear form

$$A(k) = 0 \text{ if } k \le K, \qquad A(k) = \pi(k - K) \text{ if } k \ge K,$$

where $K \ge 0$ is a specified threshold family size and $\pi \ge 0$. The Israeli child allowance formula has always been approximately of this form.

Optimal family size is

$$k_{tj} = \operatorname{argmax}_{k \ge 0} \{(a_j + 2wS_{t-1})k - (b + w)k^2 + c\pi(k - K)\cdot 1[k \ge K]\}.$$

Optimal family size varies discontinuously with parameter a_j, the discontinuity occurring at a *pivot value*

$$a(S_{t-1}, \lambda) = 2(b + w)K - 2wS_{t-1} - c\pi/2.$$

The optimal size is

$$k_{tj} = \frac{a_j + 2wS_{t-1}}{2(b + w)} \leq K \qquad \text{if} \quad a_j \leq a(S_{t-1}, \lambda)$$

$$k_{tj} = \frac{a_j + 2wS_{t-1} + c\pi}{2(b + w)} \geq K \qquad \text{if} \quad a_j \geq a(S_{t-1}, \lambda).$$

The mean family size among women who choose fertility at date t is

$$S_t = \frac{1}{2(b + w)} [\mu + 2w S_{t-1} + c\pi \cdot \text{Prob}(a \geq a(S_{t-1}, \lambda)].$$

$S(S_{t-1}, \lambda)$ is increasing and continuous in S_{t-1}; hence the equation $S_t = S(S_{t-1}, \lambda)$ generates a monotone time path for mean family size.

A steady state, where $S^o = S(S^0, \lambda)$, is defined implicitly by the condition

$$S^o = \frac{1}{2b} [\mu + c\pi \cdot \text{Prob}(a \geq a(S^o, \lambda)].$$

Solutions lie between the steady state $\mu/2b$ without allowances and the steady state $(\mu + c\pi)/2b$ with allowances proportional to family size.

Implicit differentiation reveals that a more generous child allowance formula generates increases in steady-state mean family size.

When $\pi > 0$ and $K > 0$, stronger social interactions reinforce the deviation of steady state fertility away from the threshold K. Thus, we have

<u>Proposition</u>:

Under a program with allowances proportional to family size, social interactions have no effect on mean steady state fertility S^o.

Under a program with piecewise linear allowances and threshold K, stronger social interactions reinforce the deviation of S^o from K.

Findings for Specific Distributions of Utility Parameters

The above partially characterizes steady state fertility when the allowance formula is piecewise linear, but important questions remain open.

In particular, we would like to know whether the steady state equation can have multiple solutions and we would like to understand the role of the social interaction parameter w in determining the steady state.

To partially address these questions, we examine the dynamics of fertility under two hypotheses about the shape f(a) of the distribution of utility parameters: f(a) is uniform and f(a) is log-normal.

Some configurations generate multiple steady states, some being locally stable and others not.

Performing numerical calculations in particular cases, we are able to demonstrate that small changes in the magnitude of allowances or in private fertility preferences can, in principle, yield rather large changes in mean fertility; on the order of the changes observed in Israel between 1950 and 1980.

Attempted Structural Empirical Inference

The rates of completed fertility observed in different ethnic-religious groups at different times could be produced by cross-sectional and time series variation in utility parameters, cross-sectional and time series variation in the child allowance formula, or social interactions within and across groups.

Many combinations of these forces could have generated the observed patterns.

Observation that woman j chooses family size k_{tj} implies only that the composite utility parameters $(a_j + 2\sum_m w_{mj} S_{tm}, b_j + \sum_m w_{mj}, c_j)$ satisfy the revealed preference inequalities

$$(a_j + 2\sum_m w_{mj} S_{tm})(k_{tj} - k) - (b_j + \sum_m w_{mj})(k_{tj}^{\,2} - k^2) + c_j[A_{tj}(k_{tj}) - A_{tj}(k)] \;\geq\; 0, \quad k \neq k_{tj}.$$

The identification problem is especially severe when the allowance is proportional to the number of children. Then choice of k_{tj} children is consistent with any value of the parameters (a_j, b_j, c_j; w_{mj}, $m \in M$) such that

$$k_{tj} - \tfrac{1}{2} < \frac{a_j + 2\sum_m w_{mj} S_{tm} + c_j \pi_{tj}}{2(b_j + \sum_m w_{mj})} < k_{tj} + \tfrac{1}{2}.$$

Thus, fertility preferences cannot be learned from observation of choice behavior alone.

The available data on fertility choices must be combined with a priori restrictions on the distribution of preferences.

We introduce these further assumptions:

A. Women choose family size six-to-ten years following marriage, and base their utility calculations on the average annual benefits in effect during that five-year period.

B. We define reference groups by the same criteria of ethnic origin and religiosity used in our description of fertility patterns in Israel. We also distinguish members of these groups by their year of marriage.

C. We constrain the vector (w_m, $m \in M$). Women who marry in a given year are influenced only by women who married in the five preceding years. They give equal weight w_0 to the fertility of all women married in the past five years and additional weight w_1 to members of their own group.

D. We permit the distribution of preferences to vary freely across ethnic-religious groups and year of marriage.

E. We assume all women who belong to the same ethnic-religious group and who marry in the same year to have the same values of (b, c, w_0, w_1). We permit heterogeneity only in parameter a, assumed distributed normal across women who belong to the same group and marry in the same year.

These assumptions only identify some composite parameters if the benefit is proportional to family size.

A necessary condition for identification of the structural parameters is that the benefit be nonlinear in family size.

The Israeli child allowance benefit has, to varying degrees over the years, been piecewise-linear in family size.

Nevertheless, when we attempted to estimate the model, we learned that piecewise-linearity of the formula provides an insubstantial foundation for empirical analysis.

The objective functions used in maximum likelihood or least squares estimation of the ordered-probit model were close to flat when evaluated at alternative parameter values that would formally be observationally equivalent if the allowance benefit were proportional to number of children.

Hence, under the maintained assumptions, we were unable to obtain reliable, stable estimates of the parameters describing the distribution of utility.

This empirical finding is unpleasant but instructive. It cautions against reliance on nonlinearities in child allowance formulas to identify fertility preferences.

The difficulties that we encountered are reminiscent of those that have afflicted attempts by empirical researchers to learn the effects on labor supply of benefit/tax schedules that are nonlinear in hours worked.

Social Learning from Private Experiences:
The Dynamics of the Selection Problem
(*Review of Economic Studies*, 2004)

I analyze social interactions that stem from the successive endeavors of new cohorts of heterogeneous decision makers to learn from the experiences of past cohorts.

Illustrations

Persons diagnosed with an illness may seek to evaluate the treatment alternatives by observing the experiences of persons who were previously diagnosed with the same illness and who were treated in different ways.

Youth deciding whether to initiate a risky behavior drug use may draw lessons from the experiences of their peers.

Social scientists have wanted to understand how decision makers learn about and choose innovations.

A common scenario supposes a status quo with known attributes.

A new alternative with unknown attributes becomes available.

Then successive cohorts of decision makers choose between the status quo and the innovation, with later cohorts observing the experiences of earlier ones and possibly learning from them.

The Dynamics of Social Learning

Recurrent decision problems generate dynamic processes of social learning from private experiences.

The members of each new cohort of decision makers attempt to learn from the actions chosen and outcomes realized by past cohorts, and then make decisions that produce new experiences observable by future cohorts.

The question of interest is to understand the resulting dynamic.

The Decision Maker's Inferential Problem

How may decision makers learn from the experiences of past cohorts?

This broad question has no single answer.

The inferences that a person can draw from empirical evidence necessarily depend on what he observes and on what prior information he brings to bear.

Past Research

The observational, informational, and behavioral assumptions maintained in studies such as Banerjee (1992), Bikhchandani, Hirschleifer and Welch (1992), Manski (1993a), Ellison and Fudenberg (1995), Foster and Rosenzweig (1995), McFadden and Train (1996), and Smith and Sorenson (2000) differ considerably from one another.

Some authors assume that new decision makers only observe past actions. Some assume that the outcomes of these actions are observable as well.

Some authors assume that decision makers possess private information about the outcomes associated with alternative actions. Others assume that all information about outcomes is common knowledge.

Some authors assume that decision makers can recognize earlier actors of the same *type*, who share their preferences. Others assume that unobservable heterogeneity in preferences prevents complete recognition of types.

Statistical Inference and Identification

Research has focused on the statistical aspect of social learning; that is, on inference from sample data.

Most authors assume that decision makers have prior subjective distributions over the objects to be learned and use sample data on the experiences of earlier decision makers to update these distributions by Bayes Rule.

However, Manski (1993a) assumed that decision makers apply frequentist nonparametric regression methods, and Ellison and Fudenberg (1995) posed a model of boundedly rational learning.

In contrast, this paper examines social learning when decision makers face an identification problem, the *selection problem*, as they seek to learn from observation of past actions and outcomes.

Identification Problems and Decisions Under Ambiguity

Identification problems generate ambiguity about the identity of optimal actions

A common source of ambiguity is incomplete knowledge of a probability distribution describing a relevant population – the decision maker may know only that the distribution of interest is a member of some set of distributions.

This is the generic situation of a decision maker who seeks to learn a population distribution empirically, but whose data and prior information do not suffice to identify the distribution.

Thus, identification problems in empirical analysis induce ambiguity in decision making.

The Selection Problem

Among the identification problems that confront empirical researchers and decision makers, the selection problem looms large.

The problem is that only the outcomes of chosen actions are observable; one cannot observe the outcomes that would have occurred if persons had selected other actions.

The logical impossibility of observing counterfactual outcomes poses a fundamental difficulty for empirical research in the social sciences.

It is no less an impediment to social learning.

Previous studies have assumed that new decision makers possess enough prior information to be able to infer the distribution of counterfactual outcomes.

It has been common to suppose that the selection of actions by past decision makers emulates a randomized experiment.

This solves the selection problem, provided that some past decision makers chose each feasible action(Manski, 1993a).

Note: The requirement that some past decision makers chose each feasible action is non-trivial. The phenomena of *herd behavior* and *information cascades* analyzed by Banerjee (1992) and Bikhchandani, Hirschleifer and Welch (1992) arise in part because, in their models, successive decision makers of the same type find it optimal to choose the same actions.

Maintained Assumptions

This paper examines how social learning can occur when decision makers cannot solve the selection problem.

I assume that new decision makers have no prior information about the outcomes associated with alternative actions, nor about the decision processes of earlier cohorts.

They only observe the actions chosen by earlier cohorts and the outcomes that they experienced.

I assume that new decision makers must choose actions at a specified time and cannot revise their choices once made.

Thus, they cannot undertake *learning-by-doing* and cannot otherwise wait for empirical evidence to accumulate before making decisions.

Thus, the dynamics analyzed in the paper emerge purely out of the process of social learning across successive cohorts.

Individuals do not themselves face dynamic choice problems.

Cohorts of decision makers are groups of persons who share certain observable characteristics and who make decisions at the same time. In a medical context, for example, a cohort could be a group of individuals with common demographic attributes who are newly diagnosed with a specified illness in a given year.

I assume one regularity condition and one form of prior information.

Stationarity of Outcome Distributions: The regularity condition is that, for each feasible action, successive cohorts of decision makers share the same distribution of outcomes.

Stationarity is Common Knowledge: The informational assumption is that decision makers know about this stationarity.

Findings on The Process of Information Accumulation

Suppose that decision makers want to learn the response distribution describing the outcomes that would occur if persons observationally identical to themselves would choose a given action.

Under the stationarity assumption, empirical evidence accumulates over time.

Consider the cohort of period T, who observe cohorts t = 1, . . . , T-1.

Observation of any one of these previous cohorts restricts the distribution of response to a set of feasible distributions.

Observation of all T-1 previous cohorts restricts the distribution of response to the intersection of T-1 such sets of distributions.

Thus, the class of feasible response distributions shrinks over time.

Information accumulation takes the form of *sequential reduction in ambiguity*.

Findings on Decision Making

Decision makers have bounded utility functions and aim to maximize objective expected utility.

This goal may not be achievable if response distributions are incompletely identified, but decision makers may learn that some actions are dominated.

Successive cohorts of decision makers may be able to shrink the set of undominated actions.

The time path of adoption of an innovation depends critically on how decision makers act under ambiguity.

The Process of Information Accumulation

At each date $T \geq 1$, each member of a cohort J_T of decision makers must choose an action from a finite time-invariant choice set C.

Each person $j \in J_T$ has a *response function* $y_j(\cdot): C \to Y$ that maps actions into outcomes, which take values in space Y.

Let $z_j \in C$ denote the action chosen by person j. Then person j realizes outcome $y_j \equiv y_j(z_j)$. The other outcomes $y_j(c)$, $c \neq z_j$ are latent variables.

Each cohort J_T has the structure of a probability space, say (J_T, Ω_T, P_T).

The members of cohort J_T want to learn the response distributions $\{P_T[y(c)], c \in C\}$ before choosing their actions.

Assumption 1 (Observation of Past Actions and Outcomes):

Let $T \geq 2$. Before choosing actions, the members of cohort J_T observe the distributions $\{P_t(y, z), 1 \leq t \leq T\text{-}1\}$ of actions chosen and outcomes realized by earlier cohorts.

Assumption 2 (Stationarity of Response Distributions):

For each $c \in C$, there exists a time-invariant probability distribution $P[y(c)]$ on the outcome space Y such that $P_T[y(c)] = P[y(c)], \ \forall T \geq 1$.

Observation of past experiences enables successive cohorts to draw increasingly strong conclusions about their common response distributions $P[y(c)]$, $c \in C$. The basic finding is

Proposition 1

Let Assumptions 1 and 2 hold. Let $T \geq 2$ and $c \in C$. Let Γ denote the set of all probability distributions on Y. Then the members of cohort J_T learn that

$$P[y(c)] \in H(T, c) \equiv \bigcap_{1 \leq t \leq T-1} \{P_t(y|z = c)P_t(z = c) + \gamma \cdot P_t(z \neq c),\ \gamma \in \Gamma\}.$$

Moreover, $\{P[y(c)], c \in C\} \in [H(T, c), c \in C]$. □

The proposition describes a process of *sequential reduction of ambiguity*.

Information accumulation is monotone, so there exists a *terminal information state*.

That is, there exists a $[H(c), c \in C]$ such that

$$\lim_{T \to \infty} [H(T, c), c \in C] = [H(c), c \in C].$$

Proof: For each t = 1, . . . , T - 1,

$$
\begin{aligned}
P[y(c)] &= P_t[y(c)] \\
&= P_t[y(c)|z = c]P_t(z = c) + P_t[y(c)|z \neq c]P_t(z \neq c) \\
&= P_t(y|z = c)P_t(z = c) + P_t[y(c)|z \neq c]P_t(z \neq c).
\end{aligned}
$$

Assumption 2 gives the first equality and the Law of Total Probability gives the second. The third equality holds because y(c) is the realized outcome when z = c.

By Assumption 1, cohort J_T observes the outcome distribution $P_t(y|z = c)$ and $\{P_t(z = c), P_t(z \neq c)\}$. However $P_t[y(c)|z \neq c]$ is an unobservable distribution of latent outcomes, hence an unknown element of Γ. Hence

$$P[y(c)] \in \{P_t(y|z = c)P_t(z = c) + \gamma \cdot P_t(z \neq c),\ \gamma \in \Gamma\}.$$

The feasible values of P[y(c)] satisfy this for t = 1, . . , T-1.

The second part holds because $\{P_t[y(c)|z \neq c], c \in C\}$ may be any element of the product space $\Gamma^{|C|}$. Hence the vector $\{P[y(c)], c \in C\}$ of response distributions is an element of $[H(T, c), c \in C]$.

A distribution is feasible if and only if the probability it places on each measurable subset of Y is no less than an easily computed lower bound.

If Y is countable, it suffices to consider the probability placed on each atom of Y.

Corollary 1:

Let Assumptions 1 and 2 hold. Let $T \geq 2$ and $c \in C$. Let $\eta \in \Gamma$ be a specified probability distribution on Y. Given any measurable set $A \subset Y$, define

$$\pi_{Tc}(A) \equiv \max_{1 \leq t \leq T-1} P_t(y \in A | z = c) P_t(z = c).$$

(a) Then $\eta \in H(T, c)$ if and only if $\eta(A) \geq \pi_{Tc}(A)$, $\forall A \subset Y$.

(b) Let Y be countable. Then $\eta \in H(T, c)$ if and only if $\eta(y) \geq \pi_{Tc}(y)$, $y \in Y$.

(c) Let Y be countable. Let $S(T, c) \equiv \sum_{y \in Y} \pi_{Tc}(y)$. Then $H(T, c)$ contains a unique distribution if and only if $S(T, c) = 1$. If $S(T, c) = 1$, the unique feasible distribution is $\eta_{Tc}(y) \equiv \pi_{Tc}(y)$, $y \in Y$. □

Decision Making

Let $U_j(\cdot, \cdot): C \times Y \rightarrow R^1$ denote the utility function that person j uses to evaluate actions. Thus, $U_j[c, y(c)]$ is the utility of action c to person j.

The problem is that person j does not know the outcome y(c) that he would experience under action c.

How might a person behave in this setting?

A pervasive idea in research on social learning has been that a person views himself as a member of some observable reference group and predicts that, if he were to choose a given action, he would experience an outcome drawn at random from the distribution of outcomes in this group.

Formally, assume that person j views himself as a member of cohort J_T, predicts that his outcome under each action c $\in$ C is drawn from P[y(c)], and aims to solve the problem

$$\max_{c \in C} \int U_j[c, y(c)]dP[y(c)].$$

This problem generically is not solvable under Assumptions 1 and 2, because decision makers only learn that $\{P[y(c)], c \in C\} \in [H(T, c), c \in C]$.

However, decision makers can eliminate actions that are dominated.

Assumption 3 (Elimination of Dominated Actions):

Let Assumptions 1 and 2 hold. Let $T \geq 2$ and $j \in J_T$.

For $c \in C$ and $\gamma \in \Gamma$, let $\int U_j(c, y)d\gamma$ be the expected utility of action c if outcome y were distributed γ.

Action $c' \in C$ is *dominated* if there exists an action $c'' \in C$ such that

$$\int U_j(c', y)d\eta' \leq \int U_j(c'', y)d\eta'', (\eta', \eta'') \in [H(T, c'), H(T, c'')]$$

$$\int U_j(c', y)d\eta' < \int U_j(c'', y)d\eta''$$
$$\text{for some } (\eta', \eta'') \in [H(T, c'), H(T, c'')].$$

Person j does not choose a dominated action. □

Choice Among Undominated Actions

There is no optimal way to choose an undominated action.

Hurwicz (1951) suggested maximization of a weighted average of the minimum and maximum values of the objective function that are feasible for each action. Thus, person j would solve the problem

$$\max_{c \in C} \; \lambda_j \{ \inf_{\eta \in H(T, c)} \int U_j[c, y(c)]d\eta \} + (1 - \lambda_j) \{ \sup_{\eta \in H(T, c)} \int U_j[c, y(c)]d\eta \}$$

for some $\lambda_j \in [0, 1]$.

This family of decision rules provides a simple way of expressing degrees of pessimism and optimism.

$\lambda_j = 1$ means that person j uses the maximin rule and $\lambda_j = 0$ that he uses the maximax rule.

Learning About and Choosing Innovations

Assumption 4 generates a rich variety of dynamics.

Assumption 4: Assumption 3 holds. Moreover,

(a) The choice set is C = (e, n).

At T = 1, all persons choose action e.

The outcome space is Y = {0, 1}.

The utilities that person j associates with actions e and n are

$$U_j[e, y(e)] = y(e) \text{ and } U_j[n, y(n)] = y(n) + u_j,$$

where $u_j \in R^1$. Person j knows u_j before choosing an action.

(b) Person j uses the Hurwicz criterion with parameter λ_j to choose among undominated actions.

(c) There exists a time-invariant probability distribution $P[y(e), y(n), u, \lambda]$ such that

$$P_T[y(e), y(n), u, \lambda] = P[y(e), y(n), u, \lambda], \;\; \forall\, T \geq 1.$$

The distribution of u is continuous.

Proposition 3 describes the time path of adoption of the innovation.

<u>Proposition 3</u>:

Let Assumption 4 hold. At each date $T \geq 2$,

$$P_T(z = n) =$$

$$P\{\lambda \cdot \pi_{Tn}(1) + (1 - \lambda) \cdot [1 - \pi_{Tn}(0)] + u > P[y(e) = 1]\}. \quad \square$$

Here $\pi_{2n}(0) = \pi_{2n}(1) = 0$. $\pi_{(T+1)n}(1)$ and $\pi_{(T+1)n}(0)$ are given by the updating rule

$$\pi_{(T+1)n}(1) = \max[\pi_{Tn}(1), \; P_T(y = 1 | z = n)P_T(z = n)],$$
$$\pi_{(T+1)n}(0) = \max[\pi_{Tn}(0), \; P_T(y = 0 | z = n)P_T(z = n)].$$

If $P[y(n) = 1]$ were known, Assumption 3 implies choice of n if $P[y(n) = 1] + u > P[y(e) = 1]$ and choice of e otherwise.

The rate of adoption of the innovation would be $p^* \equiv P\{P[y(n) = 1] + u > P[y(e) = 1]\}$.

However, $\pi_{Tn}(1) \leq P[y(n) = 1] \leq 1 - \pi_{Tn}(0)$. Hence the actual fraction of a cohort who choose n can be below or above p^*.

If everyone uses the maximin rule ($\lambda = 1$), the fraction of a cohort who choose the innovation weakly increases with time, but always remains less than or equal to p^*.

If all use the maximax rule ($\lambda = 0$), the fraction who choose the innovation weakly decreases with time, but always remains greater than or equal to p^*.

Conclusion

Social learning from private experiences is a process of *complexity within regularity*.

The process is complex because the dynamics of learning and the properties of the terminal information state flow from the subtle interaction of information accumulation and decision making.

Yet a basic regularity constrains how the process evolves, as accumulation of empirical evidence over time weakly reduces the ambiguity that successive cohorts face.

Theoretical analysis and computational experiments are valuable in understanding complex economic processes, but I see a pressing need for new empirical research as well.

In the present context, the critical empirical question is how decision makers actually cope with ambiguity.

The way decision makers choose among undominated actions can critically affect the dynamics of learning and choice.

An improved empirical understanding of decision making under ambiguity is necessary to guide theoretical and computational research in productive directions.