

Social Learning from Private Experiences: The Dynamics of the Selection Problem

CHARLES F. MANSKI

Northwestern University

First version received June 2001; final version accepted October 2002 (Eds.)

I analyse social interactions that stem from the successive endeavours of new cohorts of heterogeneous decision makers to learn from the experiences of past cohorts. A dynamic process of information accumulation and decision making occurs as the members of each cohort observe the experiences of earlier ones, and then make choices that yield experiences observable by future cohorts. Decision makers face the *selection problem* as they seek to learn from observation of past actions and outcomes, while not observing the counterfactual outcomes that would have occurred had other actions been chosen. Assuming that all cohorts face the same outcome distributions, I show that social learning is a process of sequential reduction in ambiguity. The specific nature of this process, and its terminal state, depend critically on how decision makers make choices under ambiguity. I use the problem of learning about innovations to illustrate.

1. INTRODUCTION

This paper analyses social interaction processes that stem from the successive endeavours of new cohorts of heterogeneous decision makers to learn from the experiences of past cohorts. Persons newly diagnosed with an illness may seek to evaluate the available 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 behaviour such as drug use may draw lessons from the experiences of their peers. These and many similarly 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.

How may decision makers learn from the experiences of past cohorts? This broad question obviously 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. For this and other reasons, it is not realistic to think that studies of social learning will ever reveal universal truths. However, such studies can illuminate how learning may occur in varied settings.¹

1. The observational, informational and behavioural 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 Sørensen (2000) differ considerably from one another. Some authors have assumed that new decision makers only observe past actions; others have assumed, as I do here, that the outcomes of these actions are observable as well. Some authors have assumed that decision makers possess private information about the outcomes associated with alternative actions; others have assumed, as I do here, that all information about outcomes is common knowledge. Some authors have assumed that decision makers can recognize earlier actors of the same *type*, who share their preferences; others have assumed, as I do here, that unobservable heterogeneity in preferences prevents complete recognition of types.

Recent economic research has focused on the statistical aspect of social learning; that is, on the problem of inference from sample data.² 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. Econometricians have long studied how identification problems limit the conclusions that can be drawn in empirical research. Decision makers attempting to interpret the experiences of past cohorts face much the same identification problems as do empirical researchers (Manski, 1993*b*). Identification problems create difficulties for decision making because they generate ambiguity about the identity of optimal actions (Manski, 2000*a*).³

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 has long been recognized to pose a fundamental difficulty for empirical research in the social sciences. It is no less an impediment to social learning.

Previous studies of social learning have “solved” the selection problem by assuming that new decision makers possess enough prior information to be able to infer the past distribution of counterfactual outcomes, despite their unobservability. It has been particularly common to suppose that the selection of actions by past decision makers emulates a randomized experiment. Then the realized distribution of outcomes among persons who actually chose a given action is the same as the counterfactual distribution of outcomes that the other decision makers would have experienced had they chosen this action. This solves the selection problem, provided only that some past decision makers chose each feasible action.⁴

Writing on the selection problem in empirical research, I have argued that credible prior information enabling solution of the problem often is unavailable; hence researchers should want to understand the possibilities for inference in the absence of such information (e.g. Manski, 1995, Chapter 2). Similarly, decision makers seeking to learn from past experiences may not possess credible prior information enabling them to solve the selection problem. This paper examines how social learning may occur when decision makers cannot solve the selection problem, but must cope with it as they can.

To engage the issue, I assume here 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

2. Most authors have assumed 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 (1993*a*) assumed that decision makers apply frequentist non-parametric regression methods, and Ellison and Fudenberg (1995) posed a model of boundedly rational learning.

3. A decision maker with a known choice set who wishes to maximize an unknown objective function is said to face a problem of decision under ambiguity. 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.

Concern with decisions under ambiguity dates back at least to Keynes (1921) and Knight (1921), who used the term *uncertainty*. Ellsberg (1961) posed the problem in a particularly evocative way through a thought experiment requiring subjects to draw a ball from either of two urns, one with a known distribution of colours and the other with an unknown distribution of colours. The formal study of principles for decision making under ambiguity dates back at least to Wald (1950), who proposed the maximin rule. Walley (1991) surveys a large body of subsequent theoretical research, while Camerer and Weber (1992) survey the limited body of empirical research on behaviour under ambiguity.

4. The requirement that some past decision makers chose each feasible action is non-trivial. The phenomena of *herd behaviour* and *information cascades* analysed by Banerjee (1992) and Bikhchandani *et al.* (1992) arise in part because, under the maintained assumptions of these authors, successive decision makers of the same type find it optimal to choose the same actions. However, the Manski (1993*a*) model of learning by decision makers with idiosyncratically heterogeneous preferences implies that a positive fraction of past decision makers do choose each feasible action.

that they experienced. Whereas research focusing on the statistical aspect of social learning has usually assumed that new decision makers observe a finite random sample of past actions and outcomes, I suppose that new decision makers observe all past actions and outcomes. This simplifying assumption is made to keep attention focused on the fundamental identification problem posed by the unobservability of counterfactual outcomes. A discussion of statistical inference appears in Section 2.2.

I assume that new decision makers must choose their 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. This simplifying assumption implies that each decision maker faces a single choice problem with predetermined information. Thus, the dynamics analysed in this paper emerge purely out of the process of social learning across successive cohorts. Individuals do not themselves face dynamic choice problems.

The *cohorts* of decision makers envisioned in this paper are groups of persons who share certain observable characteristics. 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. To simplify the analysis, I assume that decision makers do not have the latitude to view themselves as members of multiple groups; thus, the cohorts are predetermined. See Section 3 for further discussion.

This paper assumes only one regularity condition and one form of prior information. The regularity condition is that, for each feasible action, successive cohorts of decision makers share the same distribution of outcomes. The informational assumption is that decision makers know about this stationarity. Obviously, any attempt to learn from history must assume some form of continuity from the past to the present. The stationarity assumption maintained here implies that empirical evidence accumulates over time, each successive cohort being able to draw inferences from a longer history of past experiences. I show that it enables some degree of social learning to take place.

The basic finding on information accumulation, presented in Section 2, extends ideas introduced in Manski (1990, 1994), which analysed the selection problem in the presence of certain forms of exclusion restrictions. It is easy to see that the maintained stationarity assumption generates exclusion restrictions. Consider the cohort of period T , who observe the actions chosen and outcomes realized by the cohorts of $t = 1, \dots, T-1$. It is shown in Section 2 that observation of any one of these $T-1$ previous cohorts restricts the outcome distribution associated with each feasible action to a certain set of feasible distributions, the form of this set depending on which previous cohort is observed. Observation of all $T-1$ previous cohorts restricts the outcome distribution to the intersection of $T-1$ such sets of distributions. In this manner, accumulation of empirical evidence over time successively narrows the class of feasible outcome distributions, called the *identification region*. Thus social learning is a process of sequential reduction in ambiguity.⁵

Section 3 considers how decision makers may use the available empirical evidence to choose actions. I assume that decision makers aim to maximize objective expected utility, predicting that their outcomes are drawn at random from the distribution of outcomes in their cohort. Incomplete identification of outcome distributions may make it infeasible to solve this optimization problem, in which case persons face problems of decision making under ambiguity. The only widely accepted normative criterion for decision making under ambiguity is that persons should not choose dominated actions. (An action is said to be *dominated*, or *inadmissible*, if it is inferior to

5. Various authors have previously studied a diverse set of problems of learning under ambiguity; see, among others, Dempster (1967, 1968), Shafer (1976), Manski (1981), Bewley (1988), Gilboa and Schmeidler (1993) and Walley (1996). The present characterization of learning as sequential improvement in identification appears to be new.

some other action in all feasible states of nature. In this paper, the feasible states of nature are the feasible outcome distributions.) I show that social learning enables successive decision makers with given preferences to shrink the set of undominated actions and, in this sense, improve their decision making. I do not take a stand on how decision makers choose among undominated actions, but I do discuss several proposals in the literature: the maximin rule, the Hurwicz criterion and Bayes rules.

The manner in which decision makers choose among undominated actions can substantially affect the process of social learning. To illustrate, Section 4 examines the familiar problem of learning about innovations. Analysis of a simple, but hardly trivial, model shows that the time path of adoption of an innovation depends critically on how decision makers act under ambiguity. If they act pessimistically, the adoption rate of the innovation increases with time and converges to a steady state that is below the optimal rate of adoption. If they act optimistically, the adoption rate begins high and then falls to a steady state that is above the optimal rate of adoption.

Section 5 draws conclusions and suggests directions for future research.

2. THE PROCESS OF INFORMATION ACCUMULATION

2.1. Basic assumptions

To begin, I formalize the idea of a succession of cohorts who learn from past experiences. Suppose that at each integer 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 $y_j(\cdot): C \rightarrow Y$ 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 counterfactual outcomes $y_j(c)$, $c \neq z_j$ are unobservable.

To formalize needed distributional concepts, let each cohort J_T be a probability space, say (J_T, Ω_T, P_T) , with Ω_T the σ -algebra and P_T the probability measure. For each $c \in C$, let $P_T[y(c)]$ be the outcome distribution for action c in this cohort. $P_T[y(c)]$ is the outcome distribution that would be realized if a randomly drawn member of J_T were to choose c . It is not the distribution among members of J_T who actually choose c . That is $P_T[y(c) | z = c]$.

The analysis of information accumulation in this paper rests on two maintained assumptions:

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

Assumption 2 (Stationarity of Outcome Distributions). For each $c \in C$, there exists a time-invariant probability distribution $P[y(c)]$ on the outcome space Y such that $P_T[y(c)] = P[y(c)]$, $\forall T \geq 1$. This stationarity of outcome distributions is common knowledge.

Assumption 1 asserts that members of each cohort can observe the experiences of past cohorts. Assumption 2 asserts the stationarity of outcome distributions that enables decision makers of each cohort to learn about their own outcome distributions by observing past experiences. Section 2.2 shows how.

2.2. Sequential reduction of ambiguity

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

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

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

The identification region for $\{P[y(c)], c \in C\}$ is $[H(T, c), c \in C]$.

Proof. For each $t = 1, \dots, T - 1$,

$$\begin{aligned} P[y(c)] &= P_t[y(c)] = P_t[y(c) | z = c]P_t(z = c) + P_t[y(c) | z \neq c]P_t(z \neq c) \\ &= P_t(y | z = c)P_t(z = c) + P_t[y(c) | z \neq c]P_t(z \neq c). \end{aligned} \quad (2)$$

Assumption 2 gives the first equality, the law of total probability gives the second, and the third equality holds because $y(c)$ is the realized outcome when $z = c$. Inspect the right side of equation (2). By Assumption 1, cohort J_T observes the outcome distribution $P_t(y | z = c)$ and the choice probabilities $\{P_t(z = c), P_t(z \neq c)\}$. However, $P_t[y(c) | z \neq c]$ is an unobservable distribution of counterfactual outcomes, hence an unknown element of Γ . Therefore

$$P[y(c)] \in \{P_t(y | z = c)P_t(z = c) + \gamma_t \cdot P_t(z \neq c), \gamma_t \in \Gamma\}. \quad (3)$$

The feasible values of $P[y(c)]$ satisfy (3) for all $t = 1, \dots, T - 1$. Hence (1) holds.

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

Proposition 1 shows that learning is a process of sequential reduction in ambiguity. At date $T = 1$, decision makers have no knowledge at all of their outcome distributions. From $T = 2$ on, decision makers can use observations of past cohorts to learn about their outcome distributions. For each $c \in C$, the set $H(T + 1, c)$ of distributions that are feasible at date $T + 1$ is a subset of the corresponding set $H(T, c)$ at T .

Assumptions 1 and 2 form a jointly testable hypothesis. Suppose that $H(T, c)$ is empty for some $c \in C$. Then Assumption 1 or 2 must be invalid—either the observations of past experiences are in error or the outcome distributions $P_t[y(c)], 1 \leq t \leq T - 1$ vary with time. The analysis in this paper maintains Assumptions 1 and 2. Hence $[H(T, c), c \in C]$ is necessarily non-empty.

It is important to stress that Proposition 1 presumes no prior information beyond the stationarity of outcome distributions stated in Assumption 2. The proposition characterizes what the members of cohort J_T can learn from observation of past actions and outcomes if they know nothing about the decision rules used by earlier cohorts and nothing about the shapes of the outcome distributions $\{P[y(c)], c \in C\}$. If the members of J_T have further prior information, they may be able to conclude more than that $\{P[y(c)], c \in C\} \in [H(T, c), c \in C]$. For example, if members of J_T know that selection of actions at some $t < T$ emulated a randomized experiment, they can conclude that $P[y(c)] = P_t(y | z = c)$ for $c \in C$ such that $P_t(z = c) > 0$. It would clearly be of interest to determine what can be learned given other forms of prior information, but this paper does not pursue the matter.⁶

It is also important to understand that some learning is possible under conditions weaker than Assumptions 1 and 2. Suppose that the possibility of regime changes makes Assumption 2

6. To cite just one possibility, Assumption 2 could be strengthened to assert that $P_T[y(\cdot)] = P[y(\cdot)], \forall T \geq 1$, which makes the entire joint distribution $P_T[y(\cdot)] \equiv P_T[y(c), c \in C]$ stationary, not just its marginals $P_T[y(c)], c \in C$. Analysis of learning under this strengthened assumption presents a daunting challenge. The relevant research to date considers only the very special case of a two-period world with binary outcomes, and even here the analysis is complex (Balke and Pearl, 1997).

implausible. It may still be credible for members of cohort J_T to assert that their outcome distributions are the same as in the k periods prior to date T , for some $k \geq 0$. Then the conclusion to Proposition 1 holds if the intersection of sets in equation (1) is taken over the range $T - k \leq t \leq T - 1$ rather than $1 \leq t \leq T - 1$.

Or suppose that Assumption 1 is weakened to assert that new decision makers observe the experiences of finite random samples of past decision makers. Then the members of cohort J_T do not know the distributions $\{P_t(y, z), 1 \leq t \leq T - 1\}$ but can use their sample analogues, the empirical distributions of (outcomes, actions), to estimate these distributions consistently. The conclusion to Proposition 1 can be amended to read that cohort J_T can estimate $[H(T, c), c \in C]$ consistently.

To see this, suppose that, for $t < T$, a person in cohort J_T observes the actions chosen and outcomes realized by $N(t)$ randomly drawn members of cohort J_t . Let $\{P_{t,N(t)}(y, z), 1 \leq t \leq T - 1\}$ denote the empirical distributions of observed (outcomes, actions). Then the natural non-parametric estimate of $H(T, c)$ is

$$H_N(T, c) \equiv \bigcap_{1 \leq t \leq T-1} \{P_{t,N(t)}(y | z = c)P_{t,N(t)}(z = c) + \gamma_t \cdot P_{t,N(t)}(z \neq c), \gamma_t \in \Gamma\},$$

where $N \equiv [N(t), t = 1, \dots, T - 1]$. As the sample size vector $N \rightarrow \infty$, the empirical distributions $\{P_{t,N(t)}(y, z), 1 \leq t \leq T - 1\} \rightarrow \{P_t(y, z), 1 \leq t \leq T - 1\}$, almost surely. Hence $H_N(T, c) \rightarrow H(T, c)$, almost surely. For any given value of N , $H_N(T, c)$ is an intersection of sets and so inherits the monotonicity property of $H(T, c)$. That is, $H_N(T, c)$ shrinks as T increases.

The same considerations apply to Corollary 1 developed in Section 2.3 in what follows. There, the quantity $\pi_{Tc}(A)$ to be defined in equation (4) may be estimated consistently by its sample analogue.

2.3. Learning when the outcome space is countable

The characterization of the set of feasible distributions given in Proposition 1 is simple but abstract. The corollary below gives a useful alternative characterization of $H(T, c)$. The basic finding is that 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. This characterization is particularly useful when the outcome space Y is countable. Then one need only consider the probability placed on each atom of Y . This finding yields a simple quantitative measure of the size of $H(T, c)$, including a necessary and sufficient condition for existence of a unique feasible distribution.

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*

$$\pi_{Tc}(A) \equiv \max_{1 \leq t \leq T-1} P_t(y \in A | z = c)P_t(z = c). \quad (4)$$

- (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 $\eta(y) \geq \pi_{Tc}(y), \forall y \in Y$.*
- (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 $S(T, c) = 1$. When $S(T, c) = 1$, the unique feasible distribution is $\eta_{Tc}(y) \equiv \pi_{Tc}(y), y \in Y$.⁷*

7. The event $S(T, c) > 1$ cannot occur under Assumptions 1 and 2; this event implies that $H(T, c)$ is empty. The distribution η_{Tc} is distinct from π_{Tc} , which is sub-additive and hence not a probability distribution; that is, $\eta_{Tc}(y) = \pi_{Tc}(y)$ for $y \in Y$ but $\eta_{Tc}(A) \geq \pi_{Tc}(A)$ for $A \subset Y$.

Proof. (a) Suppose that $\eta \in H(T, c)$. Then, for each integer $t \in [1, T - 1]$, there exists a distribution $\gamma_t \in \Gamma$ such that $\eta = P_t(y | z = c)P_t(z = c) + \gamma_t \cdot P_t(z \neq c)$. Hence $\eta(A) \geq P_t(y \in A | z = c)P_t(z = c), \forall A \subset Y$. Hence $\eta(A) \geq \pi_{Tc}(A), \forall A \subset Y$.

Suppose that $\eta(A) \geq \pi_{Tc}(A), \forall A \subset Y$. For $1 \leq t \leq T - 1$, let $\gamma_t \equiv [\eta - P_t(y | z = c)P_t(z = c)]/P_t(z \neq c)$. Then γ_t is a probability measure and $\eta = P_t(y | z = c)P_t(z = c) + \gamma_t \cdot P_t(z \neq c)$. Hence $\eta \in H(T, c)$.

(b) Part (a) shows directly that $\eta \in H(T, c) \Rightarrow \eta(y) \geq \pi_{Tc}(y), \forall y \in Y$. Suppose that $\eta(y) \geq \pi_{Tc}(y), y \in Y$. Let $A \subset Y$. Then $\eta(A) = \sum_{y \in A} \eta(y) \geq \sum_{y \in A} \pi_{Tc}(y) \geq \pi_{Tc}(A)$. Hence $\eta \in H(T, c)$, again by (a).

(c) If $S(T, c) < 1$, the empirical evidence leaves indeterminate the allocation of $[1 - S(T, c)]$ of probability mass among the atoms of Y . Hence $H(T, c)$ contains multiple distributions. If $S(T, c) = 1$, then η_{Tc} is a probability measure. Part (b) shows that $\eta_{Tc} \in H(T, c)$.

Let η be a measure with $\eta(y) \geq \pi_{Tc}(y), y \in Y$ and $\eta(y) > \pi_{Tc}(y)$ for some $y \in Y$. Then $\sum_{y \in Y} \eta(y) > 1$, so η is not a probability measure. Hence η_{Tc} is the only element of $H(T, c)$. \parallel

The corollary shows that when Y is countable, the vector $[\pi_{Tc}(y), y \in Y]$ is a sufficient statistic for $H(T, c)$. We see immediately that $H(T, c)$ shrinks as $[\pi_{Tc}(y), y \in Y]$ increases. Moreover, we can measure the size of the region. As observed in the proof to part (c), the empirical evidence leaves indeterminate the allocation of $1 - S(T, c)$ of probability mass among the atoms of Y . It follows that the size of $H(T, c)$ as measured by the sup norm is

$$\|H(T, c)\|_{\text{sup}} = \sup_{(\eta, \eta') \in H(T, c) \times H(T, c)} \sup_{A \subset Y} |\eta(A) - \eta'(A)| = 1 - S(T, c). \tag{5}$$

The event $S(T, c) = 1$, making η_{Tc} the unique feasible value of $P[y(c)]$, occurs if $P_t(z = c) = 1$ at some $t < T$. More generally, $S(T, c) = 1$ if there exists a set of dates $\tau \subset [1, T - 1]$ such that $[P_t(y | z = c), t \in \tau]$ have disjoint support and $\sum_{t \in \tau} P_t(z = c) = 1$. Then the one feasible value of $P[y(c)]$ is $\sum_{t \in \tau} P_t(y | z = c) \cdot P_t(z = c)$. Although these cases show that $S(T, c) = 1$ can occur, it seems a rather special event. The analysis in Section 4 indicates that the generic condition when Assumptions 1 and 2 hold is $S(T, c) < 1$, implying that the available empirical evidence incompletely identifies $P[y(c)]$.

2.4. *The terminal information state*

The process of information accumulation characterized in Proposition 1 is monotone and so must converge to a *terminal information state*. That is, there necessarily exists a $[H(c), c \in C]$ such that

$$\lim_{T \rightarrow \infty} [H(T, c), c \in C] = [H(c), c \in C]. \tag{6}$$

Assumptions 1 and 2 are too weak to yield strong conclusions about the terminal information state, which depends on the interactive dynamics of information accumulation and decision making. However, these assumptions suffice to characterize the circumstances in which the terminal information state is attained by a specified date T . Corollary 2 gives this finding. Further analysis of the terminal information state will be performed in Section 4, where additional assumptions are imposed.

Corollary 2. *Let Assumptions 1 and 2 hold. Let $T \geq 2$. Let $c \in C$. Then $H(T, c) = H(c)$ if and only if $\pi_{Tc}(A) \geq \max_{t \geq T} P_t(y \in A | z = c)P_t(z = c), \forall A \subset Y$.*

Proof. By part (a) of Corollary 1,

$$\begin{aligned} \pi_{Tc}(A) \geq \max_{t \geq T} P_t(y \in A \mid z = c)P_t(z = c), A \subset Y &\Leftrightarrow \{\pi_{tc}(A) = \pi_{Tc}(A), t \geq T, A \subset Y\} \\ &\Leftrightarrow \{H(t, c) = H(T, c), t \geq T\}. \quad \parallel \end{aligned}$$

Corollary 2 takes a particularly simple form if, at each date t , the decision process of cohort J_t emulates a randomized experiment; that is, if $P_t(y \mid z = c) = P[y(c)]$, $\forall t \geq 1$. Then the corollary reduces to the statement $H(T, c) = H(c)$ if and only if $\max_{1 \leq t \leq T-1} P_t(z = c) \geq \max_{t \geq T} P_t(z = c)$.

A bit less obvious is the fact that $[H(T, c), c \in C]$ is the terminal information state if all of the distributions $P_t(y, z)$, $t \geq T$ are probability mixtures of the distributions $\{P_s(y, z), 1 \leq s \leq T-1\}$. Let $t \geq T$ and suppose that $P_t(y, z)$ is such a mixture. Then there exist non-negative numbers α_{st} , $1 \leq s \leq T-1$ such that $P_t(y, z) = \sum_{1 \leq s \leq T-1} \alpha_{st} P_s(y, z)$ and $\sum_{1 \leq s \leq T-1} \alpha_{st} = 1$. Hence, for all $A \subset Y$ and $c \in C$,

$$\begin{aligned} P_t(y \in A \mid z = c)P_t(z = c) &= P_t(y \in A, z = c) = \sum_{1 \leq s \leq T-1} \alpha_{st} P_s(y \in A, z = c) \\ &\leq \max_{1 \leq s \leq T-1} P_s(y \in A, z = c) = \pi_{Tc}(A). \end{aligned}$$

Thus $H(T, c) = H(c)$ by Corollary 2.

3. DECISION MAKING

3.1. Elimination of dominated actions

This section considers how decision makers may behave in the setting described in Section 2. For $T \geq 2$ and $j \in J_T$, let $U_j(\cdot, \cdot): C \times Y \rightarrow R^1$ denote the utility function that person j uses to evaluate actions. The utility $U_j[c, y(c)]$ that person j associates with action c may depend on the outcome $y(c)$, which the person does not know when facing the choice problem, as well as on attributes of c that the person does know.

Economists often assume that decision makers have rational expectations and choose actions that maximize expected utility. Thus person j , endowed with an information set \mathbb{I}_j , knows the objective distributions of outcomes $\{P[y(c) \mid \mathbb{I}_j], c \in C\}$ conditional on \mathbb{I}_j and solves the problem

$$\max_{c \in C} \int U_j[c, y(c)] dP[y(c) \mid \mathbb{I}_j]. \quad (7)$$

Decision makers do not have rational expectations under Assumptions 1 and 2. Person j knows only that $\{P_T[y(c)], c \in C\}$, the vector of objective distributions of outcomes within his cohort, is an element of the set $[H(T, c), c \in C]$ specified in Proposition 1. So problem (7) is not solvable, in general.

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.⁸ We can formalize this idea by assuming 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 $P_T[y(c)]$, and aims to solve the problem

$$\max_{c \in C} \int U_j[c, y(c)] dP_T[y(c)]. \quad (8)$$

8. The term *reference group* originated in social psychology with Hyman (1942). The idea that persons learn by observing the actions and outcomes of a reference group has long been prominent in social psychology and sociology, but these social scientists have remained content to theorize verbally and have not developed formal models of social learning akin to those in the recent economics literature.

Problem (8) expresses a limited-information version of the usual rational expectations model, one in which person j conditions his expectations only on the information that he belongs to cohort J_T rather than on his full information set \mathbb{I}_j .^{9,10}

Problem (8) may not be solvable if $[H(T, c), c \in C]$ contains multiple distributions. However, decision makers can eliminate actions that are dominated. The maintained behavioural premise of this paper is Assumption 3, which asserts that decision makers do not choose dominated actions:

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 another action $c'' \in C$ such that $\int U_j(c', y)d\eta' \leq \int U_j(c'', y)d\eta''$ for all $(\eta', \eta'') \in [H(T, c'), H(T, c'')]$ and $\int U_j(c', y)d\eta' < \int U_j(c'', y)d\eta''$ for some $(\eta', \eta'') \in [H(T, c'), H(T, c'')]$. Person j does not choose a dominated action.

3.2. Dominance when the outcome space is countable

The abstract characterization of dominated actions given in Assumption 3 becomes more transparent when the outcome space Y is countable and utility is bounded. In this case, Proposition 2 yields a simple description of the dominated actions.

Proposition 2. *Let Y be countable. Let $T \geq 2$ and $j \in J_T$. Let $K_{0jc} \equiv \min_{y \in Y} U_j(c, y)$ and $K_{1jc} \equiv \max_{y \in Y} U_j(c, y)$ exist for all $c \in C$. Let*

$$d \in \operatorname{argmax}_{c \in C} \sum_{y \in Y} \pi_{Tc}(y) \cdot U_j(c, y) + [1 - S(T, c)] \cdot K_{0jc}. \quad (9)$$

Action $c' \in C$ is dominated if

$$\begin{aligned} \sum_{y \in Y} \pi_{Tc'}(y) \cdot U_j(c', y) + [1 - S(T, c')] \cdot K_{1jc'} &< \sum_{y \in Y} \pi_{Td}(y) \cdot U_j(d, y) \\ &+ [1 - S(T, d)] \cdot K_{0jd} \end{aligned} \quad (10a)$$

or if

$$\begin{aligned} \sum_{y \in Y} \pi_{Tc'}(y) \cdot U_j(c', y) + [1 - S(T, c')] \cdot K_{1jc'} &= \sum_{y \in Y} \pi_{Td}(y) \cdot U_j(d, y) \\ &+ [1 - S(T, d)] \cdot K_{0jd} \end{aligned} \quad (10b)$$

$$\min[S(T, c'), S(T, d)] < 1. \quad (10c)$$

9. One may be content to view problem (8) as a conjecture about how a decision maker may behave when he does not know $\{P[y(c) | \mathbb{I}_j], c \in C\}$ but does know $\{P_T[y(c)], c \in C\}$. Yet it is natural to ask whether this decision rule has interesting normative properties. A partial normative foundation can be constructed by considering a social planner who must choose an action for each member of J_T and who wants to maximize a utilitarian social welfare function. If all members of J_T have the same utility function U , it is optimal for the planner to solve the problem $\max_{c \in C} \int U[c, y(c)]dP_T[y(c)]$. Decentralized solution of problem (8) achieves the same social welfare. If the members of J_T have heterogeneous preferences, decentralized solution of (8) continues to maximize social welfare if preferences are statistically independent of outcomes; that is, if $P_T[U(\cdot, \cdot), y(\cdot)] = P_T[U(\cdot, \cdot)]P_T[y(\cdot)]$. However, this decision rule may be sub-optimal under some forms of statistical dependence between preferences and outcomes.

10. As stated in Section 1, cohorts of decision makers are groups of persons who share certain observable characteristics. In practice, decision makers may be able to view themselves as members of multiple groups, conditioning on various subsets of the observable characteristics. It is easy to show that a social planner wanting to maximize a utilitarian social welfare function for a population of heterogeneous agents would find it optimal to define groups as narrowly as possible, conditioning on all observable characteristics (Manski, 2000a). It may be that individuals follow this prescription to specify their reference groups, but we need not assume this to be so.

Proof. Let $c \in C$. Proposition 1, Corollary 1, part (b) showed that $\eta \in H(T, c)$ if and only if $\eta(y) \geq \pi_{Tc}(y)$, $y \in Y$. Equivalently, $\eta \in H(T, c)$ if and only if $\eta(y) = \pi_{Tc}(y) + \delta(y)$, $y \in Y$, where $\{\delta(y) \geq 0, y \in Y\}$ and $\sum_{y \in Y} \delta(y) = 1 - S(T, c)$. Hence the identification region for the expected utility of c is

$$\begin{aligned} & \left\{ \int U_j[c, y(c)] d\eta, \eta \in H(T, c) \right\} \\ &= \left\{ \sum_{y \in Y} \pi_{Tc}(y) \cdot U_j(c, y) + \sum_{y \in Y} \delta(y) \cdot U_j(c, y), \delta(y) \geq 0, y \in Y; \sum_{y \in Y} \delta(y) \right. \\ &= \left. 1 - S(T, c) \right\} \\ &= \left[\sum_{y \in Y} \pi_{Tc}(y) \cdot U_j(c, y) + [1 - S(T, c)] \cdot K_{0jc}, \sum_{y \in Y} \pi_{Tc}(y) \cdot U_j(c, y) \right. \\ & \quad \left. + [1 - S(T, c)] \cdot K_{1jc} \right]. \end{aligned} \quad (11)$$

Moreover, the identification region for the vector of expected utilities for all actions is the Cartesian Product of the closed intervals on the R.H.S. of (11).

Now consider actions c' and d . By (11), the left side of inequality (10a) is the largest feasible value for the expected utility of c and the right side of (10a) is the smallest feasible value for that of d . When (10a) holds, d yields higher expected utility than c' in all feasible states of nature; that is, under all distributions in $H(T, d)$ and $H(T, c)$. When (10b) and (10c) hold, d always yields at least as high expected utility as c' and yields higher expected utility in some states of nature. Thus c' is dominated by d if either (10a) or (10b) and (10c) holds. \parallel

Proposition 2 shows that if utility is bounded, the accumulation of empirical evidence over time can enlarge the set of actions that decision makers may eliminate as dominated. Consider a sequence of decision makers who share the same bounded utility function but who make decisions at successive dates. Let $c' \in C$. The upper bound on the expected utility of this action, given on the L.H.S. of (10a), decreases with T . The greatest lower bound of the expected utility of all actions, on the R.H.S. of (10a), increases with T . Hence action c' may be undominated at early dates but dominated later on. It is not possible for c' to be dominated early but undominated later.

3.3. Choice among undominated actions

Assumption 3 leaves open how decision makers choose among undominated actions. There is no "optimal" way to make this choice, but many "reasonable" decision rules have been suggested over the years.

Wald (1950) proposed the maximin rule, which solves the problem¹¹

11. It is well known that the maximin rule is optimal in *competitive games*. In a competitive game, a decision maker who wants to maximize expected utility chooses an action from C . Then an opponent who wants to minimize expected utility chooses among the feasible states of nature. Hence solution of problem (12) yields the highest achievable expected utility if person j is a participant in a competitive game.

The maximin rule cannot be said to be optimal in individual choice problems, which lack the malicious opponent of competitive games. In individual choice problems, the appeal of the maximin rule is a personal rather than normative matter. Some decision makers may deem it essential to protect against worst-case scenarios, while others may not. Wald himself did not contend that the maximin rule is optimal, only that it is "reasonable". Considering the case in which the objective is to minimize rather than maximize an objective function, he wrote (Wald, 1950, p. 18): "a minimax solution seems, in general, to be a reasonable solution of the decision problem".

$$\max_{c \in C} \inf_{\eta \in H(T,c)} \int U_j[c, y(c)] d\eta. \quad (12)$$

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

$$\max_{c \in C} \lambda_j \left\{ \inf_{\eta \in H(T,c)} \int U_j[c, y(c)] d\eta \right\} + (1 - \lambda_j) \left\{ \sup_{\eta \in H(T,c)} \int U_j[c, y(c)] d\eta \right\} \quad (13)$$

for some $\lambda_j \in [0, 1]$. Rule (13) 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. See Gilboa and Schmeidler (1989) for an axiomatic development of the maximin rule and Hansen and Sargent (2000) for discussion of applications to macroeconomics. See Ellsberg (2001) for a discussion of the Hurwicz criterion.

Bayesian decision theorists suggest that the decision maker assert a subjective distribution on the space of feasible outcome distributions and maximize subjective expected utility with respect to this distribution. Thus person j would solve the problem

$$\max_{c \in C} \int \left\{ \int U_j[c, y(c)] d\eta \right\} dQ_{jc}(\eta), \quad (14)$$

where Q_{jc} is the subjective distribution that person j places on $H(T, c)$. It is important to understand that, despite their widespread application in economic theory, Bayes decision rules have no particular normative force in the absence of credible prior information. Berger (1985) calls attention to this when he states (page 121): “A Bayesian analysis may be ‘rational’ in the weak axiomatic sense, yet be terrible in a practical sense if an inappropriate prior distribution is used.”

The manner in which decision makers actually choose among undominated actions can substantially affect the dynamics of social learning. The broad reason, of course, is that the choices made by persons at date T generate the outcomes observed by persons at date $T + 1$, who then make their own choices based in part on this new information. Specific conclusions about the dynamic interaction of decision making and information accumulation can be drawn only if we entertain particular assumptions about choice among undominated actions. With this in mind, Section 4 studies a simple model of learning about and choosing innovations.

4. LEARNING ABOUT AND CHOOSING INNOVATIONS

Social scientists have long wanted to understand the manner in which decision makers learn about and choose innovations. A common scenario envisions an initial condition in which decision makers choose among a set of actions with known attributes. At some point, a new alternative yielding unknown outcomes becomes available. From then on, successive cohorts of decision makers choose among the expanded choice set, with later cohorts observing the experiences of earlier ones. It has often been hypothesized, and sometimes observed, that the fraction of decision makers choosing the new alternative increases with time in the manner of an S-shaped curve—first rising slowly, then rapidly, and finally converging to some limit value (*e.g.* Griliches, 1957).

Behaviour that is consistent with Assumption 3 can generate a dynamic with choice of the innovation increasing over time, but other dynamics are possible as well. The fraction of decision makers choosing the new alternative may begin high and then decrease with time, or the time path may be non-monotone. The terminal information state may convey considerable knowledge, or little knowledge, about the outcome distribution under the innovation. I show here that even a simple model can generate all of these possibilities.

4.1. *A simple model, but hardly a trivial one*

Assumption 4 specifies the model to be analysed.

Assumption 4. Assumption 3 holds. Moreover,

- (a) The choice set is $C = (e, n)$. At date $T = 1$, all persons choose action e . The outcome space is $Y = \{0, 1\}$. The utilities that person j associates with actions e and n are $U_j[e, y(e)] = y(e)$ 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 (13) with parameter λ_j to choose among undominated actions.
- (c) There exists a time-invariant probability distribution $P[y(\cdot), u, \lambda]$ such that $P_T[y(\cdot), u, \lambda] = P[y(\cdot), u, \lambda], \forall T \geq 1$. The distribution of u is continuous.

Part (a) specializes Assumption 3 in various respects. Action e is the only pre-existing alternative, which all persons choose at $T = 1$. Action n is the innovation. The outcome y is binary, taking the value zero or one. Utility functions are separable in y .

Parts (b) and (c) go beyond Assumption 3. Part (b) asserts that, to choose among undominated actions, each decision maker maximizes some weighted average of the lower and upper bounds on expected utility. Part (c) strengthens the stationarity condition asserted in Assumption 2. Successive cohorts of decision makers not only have the same outcome distributions, but have the same joint distributions of decision rules and outcomes.¹² Requiring that $P(u)$ be continuous ensures that indifference between actions e and n occurs with probability zero, so the model yields well-defined choice probabilities.

4.2. *The dynamics of choice*

Assumption 4 is simple in structure, but can generate a rich variety of dynamics. Proposition 3 describes the time path of adoption of the innovation.

Proposition 3. *Let Assumption 4 hold. At each date $T \geq 2$,*

$$P_T(z = n) = P\{\lambda \cdot \pi_{Tn}(1) + (1 - \lambda) \cdot [1 - \pi_{Tn}(0)] + u > P[y(e) = 1]\}. \quad (15)$$

Proof. Let $j \in J_T$. By part (a) of Assumption 4,

$$\int U_j[e, y(e)] dP_T[y(e)] = P[y(e) = 1] \quad \int U_j[n, y(n)] dP_T[y(n)] = P[y(n) = 1] + u_j.$$

All persons choose e at $T = 1$, so the empirical evidence identifies $P[y(e) = 1]$. The empirical evidence does not identify $P[y(n) = 1]$ but, by (11), $\int U_j[n, y(n)] dP_T[y(n)] \in [\pi_{Tn}(1) + u_j, 1 - \pi_{Tn}(0) + u_j]$. Hence, by (13),

$$\begin{aligned} \lambda_j[\pi_{Tn}(1) + u_j] + (1 - \lambda_j)[1 - \pi_{Tn}(0) + u_j] > P[y(e) = 1] &\Rightarrow z_j = n \\ \lambda_j[\pi_{Tn}(1) + u_j] + (1 - \lambda_j)[1 - \pi_{Tn}(0) + u_j] < P[y(e) = 1] &\Rightarrow z_j = e. \end{aligned}$$

Equation (15) follows from this and from the fact that u has a continuous distribution. \parallel

12. Part (c) does not strengthen the information condition asserted in Assumption 2. Here, as before, decision makers only know that outcome distributions are stationary.

Proposition 3 shows how adoption of the innovation depends on $[\pi_{Tn}(1), \pi_{Tn}(0)]$, which weakly increase with time as empirical evidence accumulates. Consider date $T = 2$. No one chose action n at $T = 1$, so $\pi_{2n}(0) = \pi_{2n}(1) = 0$. Hence $P_2(z = n) = P\{1 - \lambda + u > P[y(e) = 1]\}$. From then on, $\pi_{(T+1)n}(1)$ and $\pi_{(T+1)n}(0)$ are given by the updating rule

$$\pi_{(T+1)n}(1) = \max[\pi_{Tn}(1), P_T(y = 1 | z = n)P_T(z = n)] \tag{16a}$$

$$\pi_{(T+1)n}(0) = \max[\pi_{Tn}(0), P_T(y = 0 | z = n)P_T(z = n)]. \tag{16b}$$

Moreover,

$$P_T(y = 1 | z = n) = P\{y(n) = 1 | \lambda \cdot \pi_{Tn}(1) + (1 - \lambda)[1 - \pi_{Tn}(0)] + u > P[y(e) = 1]\}. \tag{17}$$

Taken together, equations (15)–(17) show how $P[y(\cdot), u, \lambda]$, the stationary distribution of outcomes and decision rules, determines the dynamics of learning and choice.

The optimal decision rule, which persons would use if $P[y(n) = 1]$ were known, is to choose n if $P[y(n) = 1] + u > P[y(e) = 1]$ and choose e otherwise. Hence the optimal rate of adoption of the innovation is $P\{P[y(n) = 1] + u > P[y(e) = 1]\}$. Observe that $\pi_{Tn}(1) \leq P[y(n) = 1] \leq 1 - \pi_{Tn}(0)$. Hence the actual fraction of a cohort who choose action n can be below or above this optimal rate, depending on the population distribution of (u, λ) .

It is revealing to consider two extreme cases, one in which all decision makers use the maximin rule ($\lambda = 1$) and the other in which all use the maximax rule ($\lambda = 0$). If everyone uses the maximin rule, then $P_T(z = n) = P\{\pi_{Tn}(1) + u > P[y(e) = 1]\}$. Hence the fraction of a cohort who choose the innovation weakly increases with time, but always remains less than or equal to the optimal adoption rate. If everyone uses the maximax rule, $P_T(z = n) = P\{1 - \pi_{Tn}(0) + u > P[y(e) = 1]\}$. In this case, the fraction who choose the innovation weakly decreases with time, but always remains greater than or equal to the optimal adoption rate.

In more general settings, the fraction of a cohort who adopt the innovation may or may not vary monotonically with time. This fraction may sometimes be less and sometimes be greater than the optimal adoption rate. The specific time path of $P_T(z = n)$ depends on the joint distribution $P[y(\cdot), u, \lambda]$ of outcomes and decision rules.

4.3. The terminal information state

The terminal information state is determined by the stationary distribution of outcomes and decision rules. Some sense of the possibilities is given by considering the special case in which all persons have the same value of λ , $y(n)$ is statistically independent of u , and u has support R^1 .

Let $P(\lambda = L) = 1$, where $L \in [0, 1]$. Let $b \equiv P[y(n) = 1]$. Let $p_t \equiv P_t(z = n)$, $t \geq 1$. By (16), statistical independence of $y(n)$ and u implies that for each $T \geq 2$,

$$\pi_{Tn}(1) = P[y(n) = 1] \cdot \{\max_{1 \leq t \leq T-1} P_t(z = n)\} = b \cdot (\max_{1 \leq t \leq T-1} p_t) \tag{18a}$$

$$\pi_{Tn}(0) = P[y(n) = 0] \cdot \{\max_{1 \leq t \leq T-1} P_t(z = n)\} = (1 - b) \cdot (\max_{1 \leq t \leq T-1} p_t). \tag{18b}$$

This and Proposition 3 yield

$$\begin{aligned} p_T &= P\{L \cdot b \cdot (\max_{1 \leq t \leq T-1} p_t) + (1 - L) \cdot [1 - (1 - b) \cdot (\max_{1 \leq t \leq T-1} p_t)] \\ &\quad + u > P[y(e) = 1]\} \\ &= P\{(L + b - 1) \cdot (\max_{1 \leq t \leq T-1} p_t) + 1 - L + u > P[y(e) = 1]\}. \end{aligned} \tag{19}$$

Inspection of (19) shows that the sign of $(L + b - 1)$ determines the qualitative dynamics of decision making and information accumulation. If $(L + b - 1)$ is positive, the probability of

choosing the innovation increases with time. Thus $\max_{1 \leq t \leq T-1} p_t = p_{(T-1)}$ and equation (19) reduces to¹³

$$p_T = P\{(L + b - 1) \cdot p_{(T-1)} + 1 - L + u > P[y(e) = 1]\}. \quad (20)$$

Recall that, by assumption, $p_1 = 0$. Hence equation (20) generates a monotone increasing rate of adoption whose limit as $T \rightarrow \infty$ is the value p' yielding the smallest solution to the equation

$$p' = P\{(L + b - 1) \cdot p' + 1 - L + u > P[y(e) = 1]\}. \quad (21)$$

The terminal information state is $[\pi_n(1) = b \cdot p', \pi_n(0) = (1 - b) \cdot p']$. The value p' lies in the open interval $(0, 1)$. Hence the terminal information state is informative about b but does not completely identify it.

If $(L + b - 1)$ is non-positive, the situation is entirely different. Consider dates $T = 2, 3$ and 4. The following hold:

$$p_2 = P\{1 - L + u > P[y(e) = 1]\} \quad (22a)$$

$$p_3 = P\{(L + b - 1) \cdot p_2 + 1 - L + u > P[y(e) = 1]\} \quad (22b)$$

$$p_4 = P\{(L + b - 1) \cdot p_2 + 1 - L + u > P[y(e) = 1]\}. \quad (22c)$$

Equation (22a) holds because $p_1 = 0$, (22b) because $p_2 \geq 0$, and (22c) because $p_3 \leq p_2$. Hence information accumulation ceases at $T = 3$ and the terminal information state is $[\pi_n(1) = b \cdot p_2, \pi_n(0) = (1 - b) \cdot p_2]$.

Thus the qualitative dynamics of decision making and information accumulation depend critically on how decision makers choose among undominated actions. If they act pessimistically ($L > 1 - b$), the adoption rate of the innovation increases with time. If they act optimistically ($L \leq 1 - b$), the adoption rate begins high and then immediately falls to a steady state value. However decision makers behave, social learning takes place but remains incomplete.

Of course, these are not the only types of dynamics that may occur. Decision makers may vary in their degree of pessimism/optimism. Outcomes and decision rules may be statistically dependent. If so, the process of learning about and choosing innovations may be more complex than depicted here.

5. CONCLUSION

Two broad ideas underlie the specific findings of this paper. One is that identification problems generate ambiguity in decision making. Accumulation of information that aids in identification serves to reduce ambiguity, and thus affects decision making. Given Assumptions 1 through 3, Propositions 1 and 2 show how observation of past actions and outcomes mitigates the selection problem and may enable successive cohorts of decision makers to shrink the set of undominated actions.

The other broad idea is that the dynamics of social learning stem from the interaction of information accumulation and decision making. Today's decision makers learn from the experiences of earlier ones, and then make choices that yield experiences observable by future decision makers. Proposition 3 shows that a rich variety of dynamics can emerge even in the simple setting of Assumption 4.

13. Choice dynamics of the form (20) have become familiar in research on binary choice with endogeneous social interactions. See Manski (1993c) and Brock and Durlauf (2001). The behavioural process has usually been presumed to be one of preference interactions, wherein the utility that a person associates with choice of action n depends directly on the fraction of the previous cohort who made that choice. The present derivation shows that choice dynamics of form (20) can be generated by social learning as well.

Considered as a whole, the present analysis depicts social learning from private experiences as a process of complexity within regularity. 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.

Obviously, the specific findings of this paper rest on the particular assumptions imposed. These assumptions warrant relaxation or modification in future research. The succession of cohorts may satisfy forms of stationarity that are stronger, weaker, or different from stationarity of outcome distributions. New decision makers may have prior information about the outcomes associated with alternative actions or about the decision processes of earlier cohorts. Individual decision makers may be able to engage in learning-by-doing or otherwise face dynamic choice problems. Decision makers may observe only finite samples of past actions and outcomes, and so face problems of statistical inference as well as the selection problem. Decision makers may be able to view themselves as members of multiple reference groups. It would clearly be interesting to study social learning under these and other departures from Assumptions 1 through 3.

I also see a pressing need for new empirical research on the manner in which decision makers cope with ambiguity as they attempt to learn from the experiences of others. This paper makes plain that 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 research in productive directions. For reasons discussed in Manski (1993c, 2000b), I believe that analysis of observational and experimental choice data does not, *per se*, provide an adequate foundation for empirical research on social learning. Progress will require richer data. In particular, I think that careful elicitation and interpretation of subjective data is essential if we are to understand how persons actually form expectations and make decisions.

Acknowledgements. This research was supported in part by National Science Foundation grant SES-0001436. I am grateful for comments from Lawrence Blume, Francesca Molinari, Orazio Attanasio and anonymous referees. I have benefited from the opportunity to present this work at the 2001 Santa Fe Institute Conference on the Economy as a Complex Evolving System and in seminars at Northwestern University, the University of California at Berkeley, University College London, and the University of Kentucky.

REFERENCES

- BALKE, A. and PEARL, J. (1997), "Bounds on Treatment Effects from Studies With Imperfect Compliance", *Journal of the American Statistical Association*, **92**, 1171–1177.
- BANERJEE, A. (1992), "A Simple Model of Herd Behavior", *Quarterly Journal of Economics*, **117**, 797–817.
- BERGER, J. (1985) *Statistical Decision Theory and Bayesian Analysis* (New York: Springer).
- BEWLEY, T. (1988), "Knightian Decision Theory and Econometric Inference" (Cowles Foundation Discussion Paper No. 868, Yale University).
- BIKHCHANDANI, S., HIRSCHLEIFER, D. and WELCH, I. (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades", *Journal of Political Economy*, **100**, 992–1026.
- BROCK, W. and DURLAUF, S. (2001), "Discrete Choice with Social Interactions", *Review of Economic Studies*, **68**, 235–260.
- CAMERER, C. and WEBER, M. (1992), "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity", *Journal of Risk and Uncertainty*, **5**, 325–370.
- DEMPSTER, A. (1967), "Upper and Lower Probabilities Induced by a Multivalued Mapping", *Annals of Mathematical Statistics*, **38**, 325–339.
- DEMPSTER, A. (1968), "A Generalization of Bayesian Inference", *Journal of the Royal Statistical Society, Series B*, **30**, 205–232.
- ELLISON, G. and FUDENBERG, D. (1995), "Word-of-Mouth Communication and Social Learning", *Quarterly Journal of Economics*, **110**, 93–125.
- ELLSBERG, D. (1961), "Risk, Ambiguity, and the Savage Axioms", *Quarterly Journal of Economics*, **75**, 643–669.
- ELLSBERG, D. (2001) *Risk, Ambiguity, and Decision* (New York: Garland Publishing).
- FOSTER, A. and ROSENZWEIG, M. (1995), "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture", *Journal of Political Economy*, **103**, 1176–1209.

- GILBOA, Y. and SCHMEIDLER, D. (1989), "Maxmin Expected Utility with a Non-Unique Prior", *Journal of Mathematical Economics*, **18**, 141–153.
- GILBOA, Y. and SCHMEIDLER, D. (1993), "Updating Ambiguous Beliefs", *Journal of Economic Theory*, **59**, 33–49.
- GRILICHES, Z. (1957), "Hybrid Corn: An Exploration in the Economics of Technological Change", *Econometrica*, **25**, 501–522.
- HANSEN, L. and SARGENT, T. (2000), "Wanting Robustness in Macroeconomics", Department of Economics, University of Chicago.
- HURWICZ, L. (1951), "Some Specification Problems and Applications to Econometric Models", *Econometrica*, **19**, 343–344.
- HYMAN, H. (1942), "The Psychology of Status", *Archives of Psychology*, No. 269, 5–91.
- KEYNES, J. (1921) *A Treatise on Probability* (London: Macmillan).
- KNIGHT, F. (1921) *Risk, Uncertainty, and Profit* (Boston: Houghton-Mifflin).
- MANSKI, C. (1981), "Learning and Decision Making When Subjective Probabilities Have Subjective Domains", *Annals of Statistics*, **9**, 59–65.
- MANSKI, C. (1990), "Nonparametric Bounds on Treatment Effects", *American Economic Review Papers and Proceedings*, **80**, 319–323.
- MANSKI, C. (1993a), "Dynamic Choice in Social Settings: Learning from the Experiences of Others", *Journal of Econometrics*, **58**, 121–136.
- MANSKI, C. (1993b), "Adolescent Econometricians: How Do Youth Infer the Returns to Schooling?", in C. Clotfelter and M. Rothschild (eds.) *Studies of Supply and Demand in Higher Education* (Chicago: University of Chicago Press) 43–57.
- MANSKI, C. (1993c), "Identification of Endogenous Social Effects: The Reflection Problem", *Review of Economic Studies*, **60**, 531–542.
- MANSKI, C. (1994), "The Selection Problem", in C. Sims (ed.) *Advances in Econometrics, Sixth World Congress* (Cambridge, UK: Cambridge University Press) 143–170.
- MANSKI, C. (1995) *Identification Problems in the Social Sciences* (Cambridge: Harvard University Press).
- MANSKI, C. (2000a), "Identification Problems and Decisions Under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice", *Journal of Econometrics*, **95**, 415–442.
- MANSKI, C. (2000b), "Economic Analysis of Social Interactions", *Journal of Economic Perspectives*, **14**, 115–136.
- MCFADDEN, D. and TRAIN, K. (1996), "Consumers' Evaluation of New Products: Learning from Self and Others", *Journal of Political Economy*, **104**, 683–703.
- SHAFER, G. (1976) *A Mathematical Theory of Evidence* (Princeton, NJ: Princeton University Press).
- SMITH, L. and SØRENSEN, P. (2000), "Pathological Outcomes of Observational Learning", *Econometrica*, **68**, 371–398.
- WALD, A. (1950) *Statistical Decision Functions* (New York: Wiley).
- WALLEY, P. (1991) *Statistical Reasoning with Imprecise Probabilities* (London: Chapman and Hall).
- WALLEY, P. (1996), "Inferences from Multinomial Data: Learning about a Bag of Marbles", *Journal of the Royal Statistical Society, Series B*, **58**, 3–57.