Evolved Perception and the Validation of Simulation Models

Robert Marks,
Australian Graduate School of Management,
University of New South Wales,
Sydney NSW 2052,
Australia

An earlier version of this paper was presented at
Commerce, Complexity, and Evolution,
UNSW, Sydney 2052, February, 1996.

Revised: 8 March, 1998

Correspondence to Associate Professor Robert Marks, Australian Graduate School of Management, University of New South Wales, Sydney, NSW 2052, Australia;
E-mail: `bobm@agsm.unsw.edu.au`

ABSTRACT:

We consider the interactions among sellers in mature, iterated oligopolies. In order to calibrate and validate simulation models of strategic interactions in these markets against historical data, the researcher soon confronts the issue of the curse of dimensionality of the number of possible states of the market. The paper presents a model of perception, belief, action, and outcome, and uses information theory to explore informative partitioning, which reduces the dimensionality of the model. From a learning or evolutionary viewpoint, players will have adjusted their perceptions (their partitioning) so as to end up close to their notional equilibrium partitioning. Three measures of informative partitioning are proposed.

ACKNOWLEDGEMENTS:

I wish to thank Hermann Schnabl, David Midgley, Robert Wood, and Robert Kohn for valuable suggestions. Research support was provided by the Australian Graduate School of Management, the Australian Research Council, and the Max Planck Institute, Jena.

KEYWORDS:

repeated games, oligopoly, perceived partition, entropy, information, validation.

1. Introduction

Economics has been a predominantly analytical discipline in its search for necessity — reducing the complex phenomena of interacting agents to interactions of one or two causes, holding all else unchanging ("ceteris paribus") and then attempting to reintroduce the complexity through additivity, making the assumption of linearity. But sometimes all the king's horses and all the king's men ... and non-linearities and other complexities have stymied the search for necessity through additivity. This has led to two sorts of solutions: non-linear dynamics, in which second-order interactions are specifically modelled, and simulations, in which the search for necessary conditions is abandoned and, rather, sufficiency is the goal: what are the consequences in the aggregate of individual agents behaving just so — what assumptions at the micro level are sufficient for the emergence of a specific pattern of economic phenomena? There are many instances of the search for sufficient conditions elsewhere in this volume.

But one aspect of this research approach has been relatively neglected: the issue of validation from historical data. In the standard analytical approach, linear regression has been used to derive coefficients in the linear relationships among a dependent variable and a set of independent (causative) variables. In general, there has been no shortage of data. With the recent approach of simulation for sufficiency, there has been little use of historical data: data series for complex economic repeated interactions are uncommon, and anyway the nature of sufficiency — "*these* coefficients in *this* model may result in *this* emergent behaviour" — may not be generalisable. But there is a need for greater use of historical data, both to aid in the derivation of coefficients and to help validate the models. This latter function is important in helping convince, first, the profession and, later, skeptical policymakers that the simulation models are something more than toy stories of how the economy just might work.

A general characteristic of simulation models, however, results in difficulties when one attempts to fit them to historical data: in general, the historical agents will have had many degrees of freedom in which to act, and have used them. (Indeed, had they not done so, the models would have been very simple and the level of potential surprise (Lave & March 1975) much less, and the value of the modelling exercise much reduced.)

With continuous models, this is not an issue, but with discrete models there is a trade-off: more degrees of freedom is more realistic, but more degrees of freedom results in the phenomenon dubbed by Bellman the "curse of dimensionality" (Bertsekas & Tsitsiklis 1996). What we mean is that with many degrees of freedom, the number of potential states grows quickly, as we demonstrate below.

Surely the answer is to use continuous models and to avoid discrete models? Well, if the phenomena being modelled are such that continuity holds, so that functional relationships are continuous, then, of course, avoid discrete models. For relationships in which the variables can be classified or measured using "interval" or "ratio" scales, continuity holds, but for "nominal" scales it does not, and for "ordinal" scales it may not.

2. An Example: The Rivalrous Dance

Two petrol stations face each other across an intersection. They sell different brands of petrol, regarded as perfect substitutes by almost all motorists. As a consequence, almost all sales go to the station selling cheaper petrol at any time. Prices are not stable, but move up and down, often in an asymmetric pattern of gradual falls followed by sharp rises. (Slade 1992 describes such situations in the Vancouver market.) Sometimes, a move as small as 0.1¢ per litre by one seller will elicit a responding price change by its rival; at other times there will be no response. How can we model this interaction, while paying attention to the observations that both sellers are interested in maximising their profits over time, and that both sellers are economising on their use of information? To generalise, information of one's rival's strategic behaviour is costly to obtain and to process, so profit maximisation includes cost-minimising behaviour on the perception and use of information.

In a competitive market, that is, when there are many rivals selling perfect substitutes to a large number of buyers, this is not an issue, since firms will be selling at the going price. But when the products are not homogeneous and when the number of sellers is small — the case, for instance, of branded goods sold in oligopolies — there will be a range of prices (and other marketing instruments' levels) at any instant, and there will be a jockeying through time as firms move along their sales–profits trade-offs: sometimes going for sales and market share at the expense of profits (at lower prices) and at other times the opposite. We have analysed this rivalrous dance in the framework of strategic game theory (Midgley et al. 1997, Marks et al. 1997, Marks 1998 forthcoming).

We use historical data of actual responses by sellers competing through time in order to endogenise three aspects of a stimulus-response model of a seller in such a market. First, the way in which each seller partitions the signals he or she receives of the actions of others. Second, the way in which each seller decides on a set of possible responses. Third, the best mapping from perceived signal to best action. The metric used is average profits over many periods, where we use established market models to map from the players' actions, taken together, and each player's period profit. When completed, this research programme will allow use of historical data to estimate each seller's endogenous decisions of perception, response, and action.

In a static setting, it is relatively easy to determine whether a small price change by a rival should be responded to. In a dynamic interaction, however, the response is not obvious, since any single change may be part of a longer-term pattern or may be the start of an attempt to signal the end of a price war (if upwards, after a period of low prices) or to steal rivals' sales (if downwards, after a period of high prices). Just which it might be is not always clear, especially with costly observing, processing, and recalling of prices, present and past.

For our stimulus–response agent, the stimulus is the state of the market and the response is the combination of price and other marketing actions chosen in the next period — we confine ourselves to discrete time, since our data show only once-weekly changes in actions: an iterated oligopoly. The state of the market includes the actions of all players in previous periods, and the number of weeks remembered is endogenous. This seems to be a simple model, but when there are more than two players, or more than a handful of possible prices, or more than one period's memory of past prices, the number of possible states grows quickly. In Midgley et al. (1997), we show that if there are p players, a possible actions per period, and m periods of memory, then the number of states is given by $a^{m \times p}$. With four players, eight possible actions, and two-period memory, the number of possible states is thus almost 4.3 billion. This "curse of dimensionality" poses a severe problem for estimation of such models.

Although agents may choose their actions from a continuous set, we believe that partitioning to reduce the number of states is appropriate for two reasons. First, on grounds of tractability, and, second, since partitioning seems a good description of actual behaviour.[1] Agents apparently partition the space of possible actions and respond only when a rival's action changes the state, suitably defined using the partitioning. This behaviour is consistent with the observation that information gathering and processing is a costly activity, to be economised: small changes in a rival's price may not be worth responding to, which brings us back to the observation above: only when a rival's price change is sufficiently large to change the state will the seller (the economic actor) respond.

On the face of it, the best partioning should be that that loses the least information. We have operationalised this in Marks (1998 forthcoming), by searching for the partioning that minimises the entropy, as we discuss below. But information is not an end in itself, as we have implicitly assumed. Rather, information is a means to an end — to maximise net return in the repeated interaction among the players.[2]

The next stage of our on-going research program is to search for the best combination of information partitioning and the consequent mapping from state to action, using exogenous actions, as we have done previously. The last stage of the program must be to endogenise the choice of actions too, so that we are searching for the best combination of perceived states (partitions), action mappings, and final actions in the repeated interactions.

Such a model is discrete: since the numbers of the states of the model is arbitrary, and since we have no wish to constrain our machine-learning algorithm, the model scales are nominal: a change in input from a state, say number 24, to another, say number 25, will not necessarily result in a change of action from number 3, say, to number 4. We must use discrete formulations instead of continuous functions.

1. Think of the large number of words the Inuit have to describe snow as a measure of the fineness of their perception partition for kinds of snow.
2. As Radner (1972, p.8) puts it, "One information function [or partitioning] is better than another if the maximum expected utility achievable with the first is greater than the maximum expected utility achievable with the second." Here, this corresponds to the maximum average profit of a brand in repeated interaction with others.

This causes no difficulty for digital computing; indeed, it is the continuous models which in principle are ill-suited to digital computing. The problems with discrete models arise when we consider the number of possible states — the curse of dimensionality — as discussed above.

2.1 Coffee Sellers

In the course of our earlier work on the U.S. retail ground coffee market, in which we modelled players as responding simply to the past prices and other marketing actions of their strategic rivals (Marks et al. 1995; Midgley et al. 1997), we became aware of the importance of modelling not just the patterns of response of the strategic players, but also their perceptions, both in looking back and in discerning whether small price movements of their rivals' are strategically significant. How might such perceptions be endogenised? A firm answer to the question of how players partition[3] their perceptions of others' actions, both through time and across the price space, will also provide information on how much or how little information they choose to use: in short, how boundedly rational players are (Rubinstein 1998).

In the coffee market we observed the price to vary from about $1.50 per pound to about $3 per pound. Since cluster analysis shows that some prices and marketing actions are used more frequently than are others for each brand, the earlier studies used four of these by each brand as the actions in its simulation. But even with this severe partitioning of action space, we found that the historical profit performances could be improved by our simple four-action, one-round-memory artificial agents.

But cluster analysis is a crude technique. We wish to use the data to examine the price partitions that the players actually used. Such partitions will generally be in terms of price (and marketing action) levels, but the boundaries introduced mean that (away from the boundary) a one-cent-a-pound change in price is not a signal responded to by the other players, while that (at the boundary) such a small price change will be seen as strategically significant by the rival players. It may be that we should partition the first differences of prices, so that a small change in price will not be perceived as a strategically significant shift (no matter where the price was before the shift); only a price change (positive or negative, symmetrically?) will be seen as such.

3. Formalities

We first formalise the process that each player uses in deciding his or her action in the market from one week to the next, using a framework outlined by Lipman (1995). Each week, faced with the actual external state (or E-state), the player perceives an internal state (or P-state), which will update his beliefs, on which are conditioned his action for the week, which together

3. The concept of partitioning in order to use the coarsest (or minimal) partition which is as informative as the non-partitioned space was introduced by Blackwell & Girshick (1954).

with the actions of his strategic rivals determines his profit that week.

There is a finite set of *external* states (or E-states), Ω. The E-states are defined by the prices (and marketing actions) that each of the players determined for its brand for a large number of weeks into the past. $\Omega = A_1 \times A_2 \times A_3$, where A_i is the vector of brand i's prices (or actions) for all weeks into the past. Figure 1 illustrates the model from external E-state to final payoffs.

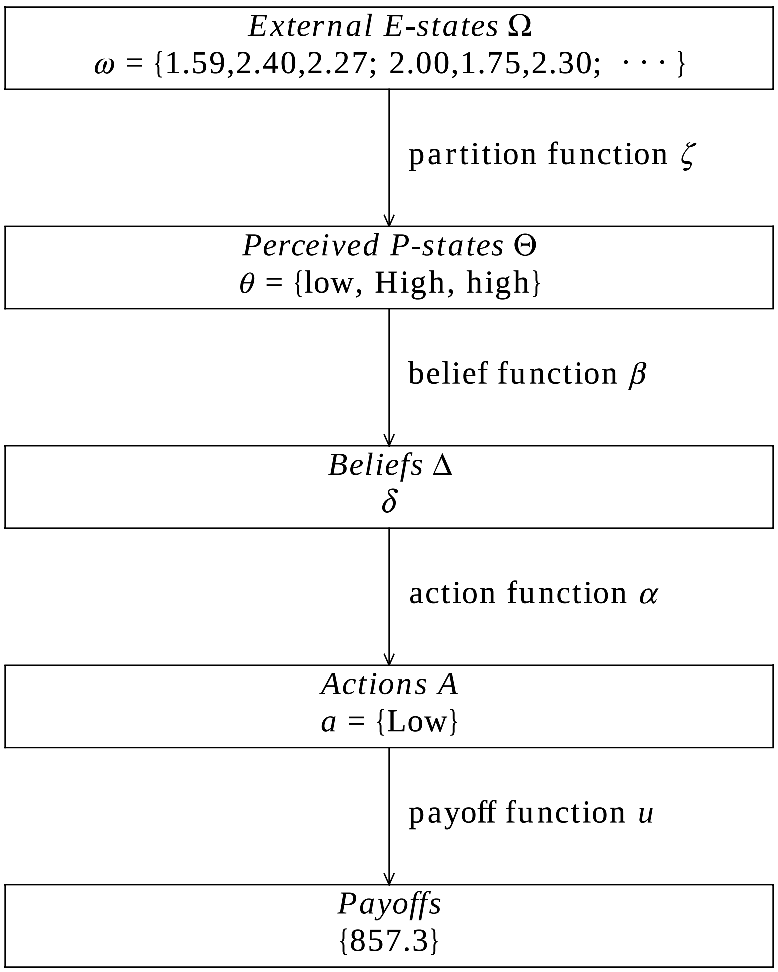

Figure 1: From External State to Payoff: The Player Modelled

But it is unlikely that players perceive the information partition at its objective fineness, as defined in the E-state. Nor is it likely that players remember more than a few weeks past in determining the internal state θ. There is a function $\zeta: \Omega \to \Theta$, that tells which *perceived* P-state θ the player observes as a function of the E-state, where ζ is the perception function: in E-state ω, the player observes P-state $\theta = \zeta(\omega)$.

As a consequence, the true information content of the P-state θ is that ω is one of the E-states generating this P-state: the true E-state ω is some element of $\zeta^{-1}(\theta)$. If the P-state is optimally determined, then the lost information is valueless to the player — he is no worse off with the coarser partition of the P-state than with the finer partition of the E-state. But if the P-state is sub-optimal, then the lost information is valuable, in that its use would result in a perception of the rivals' behaviour that would on average result in a higher profit for the player.

There will be a set of actions the player can choose from, denoted by A, with at least two elements. How or whether these actions are related to the perceived P-states is an empirical issue. Note that since such perceptions are subjective, there is no guarantee that different players will perceive the same sets of P-states.

There will also be a profit function (usually in the form of a payoff matrix): $u : A \times \Omega \rightarrow \mathrm{R}$, which describes how the state affects the value of the different actions available to any player. Note that the profit function includes the true E-state of the market during the present week (including all players' current actions), which will not be available to the players until after they have each chosen their actions. Note, too, that players will only know their perceived states, not the true external states, even later. In general, one can assume a prior probability distribution q on Ω, although in this model the probability distribution over external states is determined endogenously by the choices of the players in the market.

How does the P-state θ determine beliefs about the external state? Let Δ denote the set of probability distributions on Ω. Then β: $\Theta \rightarrow \Delta$ is the belief function. Beliefs matter because actions are contingent on them. The mapping from belief to action α: $\Delta \rightarrow A$ is the action function.

Following Lipman, we can check the internal consistency of the belief function β. No processing: $\delta = \beta(\theta)\ \forall\ \theta$, which suggests $\delta = q$, the prior distribution. Full processing: $\theta \neq \theta' \Rightarrow \beta(\theta) \neq \beta(\theta')$. We expect $\beta(\theta)$ to put probability 1 on the set $\zeta^{-1}(\theta)$. As Lipman puts it, the player should be able to say to himself: "My beliefs are δ. But I know I'd have these beliefs if and only if $\omega \in W$. So I shouldn't be putting any probability on states outside W."

Lipman distinguishes between interim optimality and ex-ante optimality. For the former, an action function α: $\alpha(\delta)$ which maximises the following function for all δ must be derived:

$$\sum_{\omega \in \Omega} u(a, \omega)\ \delta \tag{1}$$

Then a behaviour rule must be constructed by letting $f(\omega)$ equal the action $\alpha(\delta)$ where P-state $\zeta(\omega)$ results in action δ. That is, for each ω: $\beta(\zeta(\omega)) = \delta$, let $f(\omega) = \alpha(\delta)$. This describes how the player will behave in searching any given solution.

If $\beta(\zeta(\omega))$ = $\beta(\zeta(\omega'))$, then $f(\omega) = f(\omega')$. If the player has the same beliefs in two E-states, then his behaviour is the same in those E-states; that is, f is measurable with respect to $\beta(\zeta)$).

In our earlier studies, we derived equation (1), and we used firms' profits to proxy for u. Our action function was $\alpha(\delta)$, where δ is the player's belief of the E-state: we explicitly separated the determination of δ and the determination of α.

Lipman raises the question: Is it odd to model bounded rationality by assuming optimal information processing? Why not just choose optimally a given ω? Well, we assume general knowledge, that is, how to solve, not the specific solution. The model shows how to choose β and f contingent on ω. Moreover, if players do not achieve optimal β and f, then the model of the world as the player sees it is not completely specified.

3.1 Radner's Framework

Radner (1972, pp. 3–8) presents a non-strategic antecedent of our model. An act is a function from the set S of states of the world to the set C of consequences. For any act a and any state s, let $a(s)$ denote the corresponding consequence that follows from a choice of a and the occurrence of s. A function u on the set of consequences C is the utility function. A function ϕ on the set of states of the world S is the subjective probability function, such that the expected utility function U, defined on acts by

$$U(a) = \sum_s u\,[a(s)]\,\phi(s),$$

represents the ordering. (The state of the world could include others' actions, but Radner is not interested explicitly in strategic interactions.)

Let Y denote the set of alternative signals that the decision maker can receive. (These are our partitions of P-states.) The information function η associates to each state s a signal $y = \eta(s)$. (This is our partition function.) Let D be the set of alternative decisions available. Let $\rho(s,d)$ be the consequence of decision d if state s obtains. The decision is chosen according to a decision function δ, so that, if state s obtains, then the signal will be $\eta(s)$, and the decision taken will be $\delta(\eta(s))$, and the consequence will be $\rho(s,\delta(\eta(s)))$. Therefore to each information function η and each decision function δ there corresponds an act $a(s) = \rho(s, \delta(\eta(s)))$. The set of acts available depends on the set of available information and the decision functions, which are mappings from P-state to actions, in our terms..

For a given information (or partition) function η, therefore, the problem is: choose an optimal decision function δ from the set of all possible decision functions: from the set Y of signals to the set D of signals. For each signal, an optimal decision maximises the conditional expected utility of the consequence, given the signal, which is Radner's principle of maximum expected utility.

For any state s and decision d, he defines the payoff $w(s,d)$ as

$$w(s,d) = u\,[\rho(s,d)],$$

where w is the payoff function. The expected utility of an information function η and a decision function δ can be expressed as

$$U(\eta,\delta) = \sum_s \phi(s)\,w(s, \delta[\eta(s)]).$$

For the given information function, to each signal η is associated the set S_y of all states that give rise to the signal y, that is, the set of states s such that $\eta(s) = y$.

For any decision function δ that uses the information function η, all states in the same set S_y must lead to the same decision.

$\therefore$ the expected utility of (η, δ) is

$$U(\eta,\delta) = \sum_s \phi(s)\, w(s, \delta(\eta(s)))$$
$$= \sum_y P(y) \sum_{s \in S_y} P(s \mid y)\, w(s, \delta(y)),$$

where $P(y)$ is the probability of the signal y, and $P(s \mid y)$ is the conditional probability of the state s, given the signal y, that is,

$$P(y) = \sum_{s \in S_y} \phi(s)$$
$$P(s \mid y) = \phi(s)/P(y)$$

for s in S_y. We assume that $P(y) \neq 0$ for every signal y in Y.

So maximising $U(y, \delta)$ is equivalent to choosing for each signal (of a partition) y a decision d that maximises the conditional expectation

$$\sum_{s \in S_y} P(s \mid y)\, w(s, d)$$

If two decision functions use the same information (or partition) function, then we can say that the first is better than the second if (with the given information function) it gives a higher expected utility than the second. But the comparison of information functions is not as simple: we must compare the expected utility of information functions when used with their corresponding *optimal* decision functions.

3.2 Comparisons of Information Structures

McGuire (1972) presents a version of Blackwell's Theorem and discusses its importance for deriving measures of informativeness.

Blackwell's Theorem: An information structure P is regarded as "generally more informative" than information structure Q ($P \supseteq Q$) if for all payoff matrices U the set $u(Q, U)$ is contained in the set $u(P, U)$.

Let partitions or information structures P and Q possess a common (finite) state-of-the-world set S and finite signal sets Y and Z respectively. Then the following is true:

$$P \supseteq Q,$$

which imposes a very incomplete ordering on the set of all information structures or partitions.

Any search for a one-dimensional measure of "informativeness" is a vain one: there exists no real-valued function f on the set of information structures or partitions such that

$$f(P) \geq f(Q) \text{ iff } P \supseteq Q.$$

In particular, entropy cannot individually serve as an indicator of informativeness.

4. Partition Models

Information processing can be summarised by a partition Π of the set of E-states Ω. A partition Π of a set Ω is a collection of subsets of Ω with the property that every $\omega \in \Omega$ is in exactly one of these subsets. The elements of the partition Π are often referred to as *events*. Intuitively, a partition Π is said to be *finer* than a partition Π' when learning which event of Π contains a given ω conveys more information than learning only which event of Ω' contains ω; the converse is a *coarser* partition.

The partition Π is easily interpreted in terms of information processing: if Π has only one event (the entire set Ω), then the player is not processing his input P-states at all, which corresponds to the case where $\beta(\theta) = q$ for every P-state θ. By contrast, a partition that has a different event for each different E-state ω involves complete processing: the player processes the information so thoroughly that he recognises every possible distinction between inputs. His partition could not be finer.

Because Π summarises information processing, write $V(\Pi)$ instead of $V(\beta)$, for the expected profit associated with information processing according to the belief function β, which is identical to the expected profit associated with the information partition Π. If we further assume that the cost of a given information processing function β depends only on the partition β generates, then we can work with $c(\Pi)$ instead of $c(\beta)$ for the expected information processing costs.

5. Players as Stimulus-Response Machines

One reason for studying game-playing machines is that they can be used to give a formal description of the concept of "bounded rationality" (Simon 1972, Rubinstein 1998), since finite machines must by definition be bounded. concept. An automaton consists of a number of internal states, one of which is designated the initial state; a transition function, which specifies how the automaton changes states in response to the other players' actions; and an output function, which maps state to action. See Marks (1992) for a fuller treatment. In our earlier studies, (Marks et al. 1995; Midgley et al. 1997) we used the genetic algorithm to determine the initial state and the mapping from state to action.

Let I denote the set of possible histories of play (of actions). Then with three players $I = A_1 \times A_2 \times A_3$, where A_i is the history of player i's actions in the game. A strategy in the game is a function σ that specifies an action as a function of the state of the game, which in turn is a function of the history of the game. If the game has an unlimited number of rounds, then after any history h the remaining game is still infinite. Hence a strategy for the overall game, σ, specifies a continuation strategy following h for the game. Kalai & Stanford (1988) call this the induced strategy, $\sigma \mid h$. We can say that two histories, h and h', are equivalent under σ if they lead to the same induced strategy; $\sigma \mid h = \sigma \mid h'$.

Lipman argues that it is easy to show that this is an equivalence relation, so that it generates a partition of the history set I, which can be denoted by $I(\sigma)$. If the player knows which event of this partition a history lies in, then he knows enough about the history to determine the strategy it induces. Kalai & Stanford show that the number of internal states of the smallest automaton which plays a given strategy is equal to the number of sets in this partition, when the "Moore machine" representation is used.

In our earlier studies, the set of external states Ω is the set of histories $I = A_1 \times A_2 \times A_3$, where we model the strategic interaction of three brand managers as players. We arbitrarily chose a time partition of one-round memory, so that no actions of more than a week ago were directly perceived by the players (although indirect influences through others' actions last week were not, of course, excluded). To partition the large number of possible prices, we used cluster analysis on historical data of the oligopoly in order to partition the price space into four bands, again an arbitrarily chosen number. The partitions varied with brand.

These techniques allowed us to map the E-state of brands' prices (and other marketing actions) for many weeks into a much coarser P-state of one week's data, suitably partitioned; an exogenous perception function $\zeta: \Omega \rightarrow \Theta$. As described, we then used machine-learning to search for better mappings from P-state to action, or $\alpha(\beta(\theta))$. Note that, using machine representations, we did not explicitly model beliefs δ, or a belief function ($\beta: \Theta \rightarrow \Delta$), or how actions are mapped from beliefs ($\alpha: \Delta \rightarrow A$). Instead, we defined our response function as a mapping from P-state to action: $\gamma: \Theta \rightarrow A$.

The set of actions, A, is the set of strategies for the repeated game. Hence, following Lipman, any strategy σ can be described as a behaviour rule f from $I(\sigma)$ into A, where $f(h) = \sigma \mid h$. Thus we can separate the choice of a strategy σ into the choice, first, of a partition on the set of histories Π, and, second, of a function from Π to the set of strategies or actions. If used, the cost function c is usually taken as an increasing function of the number of events of the partition only, $c(\Pi)$, although other functions are possible.

6. Optimal Partitioning

Lipman (1995) discusses a class of models in which, although the E-state is observed directly, it is classified according to which of two sets it falls: whether or not it is above a certain real-valued threshold.[4] There seems no reason why the concept should not be generalised to multiple thresholds. The exogenous partitioning of our earlier studies was into four regions, requiring three thresholds, but we have considered finer partitions, although the programming effort increases in the fineness.

6.1 A First Cut

We start by considering the simplest partition of the price space, into two regions, a *dichotomous partition* between "low" and "high" prices. The question is where best to draw the boundary between the two regions. To

4. If the E-state is not already expressed as a real number, it must first be translated into a real number. In our case, however, prices are real numbers, up to the integers.

explore this issue, we set up a model in which the choice of where to divide the region between the lowest price and the highest price is one of eight points, dividing the price space into nine equal regions. See Marks (1998) for an operationalisation of these techniques.

From above, the set of external states Ω of the market with three strategic players is the set of histories $I = A_1 \times A_2 \times A_3$, but we wish to define a new set of market states based on the perceived states Θ. Instead of the set of E-state histories I, define a set of histories $\hat{I}_i = \hat{A}_{1i} \times \hat{A}_{2i} \times \hat{A}_{3i}$, where $\hat{A}_{ji}$ is the history of actions of player j as perceived by player i. As soon as we introduce subjective perceptions into the game, we introduce the possibility of subjective histories, too, but, so long as the partitioning which gives rise to the perceived actions of self and others is endogenous, no player could improve his or her payoffs by changing his or her partitioning of the price space, at least in equilibrium. From a learning or evolutionary viewpoint, players will adjust their perceptions (their partitioning) so as to end up close to their notional equilibrium partitioning.

6.2 Measures of Optimality

Which partitioning is best? To attempt to answer this question, consider the simplest non-trivial partition: high or low.

A dichotomous partition divides the price line into two regions only: "low" (below some partition point λ) and high" (above it); there remains the empirical issue of the optimal location of the dichotomous partition point. Since there is only one degree of freedom in its choice, we can plot any measure against its location. Two measures are:

1. The number of perceived states; and
2. The closely related measure of sample alog entropy across all perceived states, from equation (3).

The two measures are brand- or player-independent, since they don't require consideration of the actions that result from the perceived states, by player.

Consider the partition that loses the least amount of information.[5] One candidate is the partition (or partitions) which result in the highest number of perceived states, but there is a more informative measure: Theil (1981), in discussing the general issue of information measures associated with events, suggests entropy.[6] Entropy H is given by

5. Later we shall heed McGuire and Radner and consider a metric not of information but of profit, as a function of partitioning.
6. Of course, information is merely the means to an end: the player's profits, or expected profits in a stochastic game. But, as McGuire (1972) argues, the search for a one-dimensional measure of "informativeness" — the value of a "information structure" or partition — is in vain; entropy included. See Section 3.2 above. See also Radner (1987, p.300): "...there is no numerical measure of quantity of information that can rank all information structures [partitions] in order of value, independent of the decision problem in which the information is used".

$$H \equiv -\sum_{i=0}^{N-1} p_i \log_b p_i, \tag{2}$$

where there are N perceived states, and the probability (or observed frequency) of state i is p_i. Theil argues on axiomatic grounds that entropy is justified.

The maximum number of perceived states is equivalent to entropy as an information measure only when each state is equally likely or frequent, as is readily seen in equation (2) with $p_i = 1/N, \forall\, i$. With non-uniform distribution of states, the measure of the maximum number of states N throws away information about each state's frequency. None the less the two measures are empirically close at determining the optimal partition point with dichotomous partitioning. In order to better compare the two measures, we use the antilogarithm of entropy, or *alog entropy* (AE), which is given by the expression:

$$AE \equiv \text{antilog}_b H \equiv b^H = \frac{1}{\prod_{i=0}^{N-1} p_i^{p_i}}, \tag{3}$$

where b is the base of the logarithm used in equation (2). This measure, unlike entropy, has the additional benefit of being independent of the base b. The units of the measure of alog entropy are "equivalent states".

Marks (1998) reports empirical studies of these two measures using a data set of historical interactions in a mature, iterated oligopoly.

6.3 A Player-Specific Measure

There is the possibility of a third, brand- or player-specific measure, since there is no constraint on players to repond identically to the same observed state, and they do not.

With one-week memory, three players, and dichotomous partitioning, there are $2^3 = 8$ possible states of the market, as defined by the partition point between "low" and "high". Modelling the players as stimulus-response automata, capable of perceiving eight market states for the previous week's prices, we classify each player's prices into eight possible regions, equally spread between that brand's minimum and maximum prices.

If we have specified the model correctly, and if the partition point is optimally chosen, we might expect that there is a one-to-one mapping of perceived state to action. (With, say, four price regions, we might expect a two-to-one mapping, in which perceived information is abandoned in the choice of action.)

The brand- or player-specific measure of the mean number of action mappings per observed state given the maximum number of observed states could result in three thresholds or points of partitioning per player, π_{ij}, player i's point of partitioning player j's price actions into "low" and "high" regions. This allows subjective partitioning across players, and also each player to customise his perceptions of each of his rivals and himself.

The measure of "best" used with this measure means that the definition of states which follows from the partitioning is better when, in the limit, each action is supported by a unique state. If, in the limit, we find that we cannot reduce the number of actions per state to one across all states, then this may be thought of as one or more of several possibilities: a misspecification of the model (it may be, for instance, that players respond to the first differences in prices — price *changes* — rather than where in price space a rival player has chosen to act), or that the assumption of a deterministic state $\rightarrow$ action behaviour for each player is wrong, with some mixed-strategy element instead. Another possibility is that players did not have the data that we are using when they made their choice; and a further possibility is that weekly marketing actions (prices etc.) were decided not on a weekly basis, but beforehand over a block of weeks.

1. For a given partition point and for a given brand, determine into which region of 8 equally spaced regions between the minimum and maximum prices for that brand the next week's price falls.

2. With the discrete states for a given partition point and the price regions for a given brand as determined above, now calculate an 8×8 matrix whose elements indicate the number of times state i resulted in a price for a given brand in range j after one week, across all weeks of data. With one-week memory, 8 price regions provides the coarse, dichotomous partitioning of the price space with something to support: with fewer than 8 price ranges there is unnecessary information, which will in general be costly to obtain. The matrix is brand- or player-specific. Since the row dimension of the matrix is equal to the number of possible states, a different definition of state may result in a different number of rows for this matrix.

3. This matrix allows easy counting of the number of distinct action mappings from state to price region, for the brand or player under examination, and for the given partition point. (Some states may not appear in the data.)

4. Use a lexicographic ranking: if the number of states is less than all, then ignore this partition. If the number of states is equal to the maximum, then the best partion is that which minimises the mean number of mappings from state to price region. A (minimum) single mapping for each state would correspond to an ideal fit between partition (and states) and the brand's pricing behaviour (as segmented into price regions). Call this measure the *state-mapping brand-specific measure*.

We have not derived comparative empirical results for this third measure.

7. Conclusion

As the power of computers has grown and numerical techniques have improved, there has been a growing demand for simulation techniques in economics. This has been reflected by the emergence of new journals and conferences which specialise in the theory, development, and application of simulation techniques in economics. A relatively undeveloped area of application has been the use of simulation with historical data, and we have argued that there are particular issues and problems associated with the validation of these models, in particular the issue of the need to partition the historical data. From a learning or evolutionary viewpoint, we assume that players will have adjusted their perceptions (their partitioning) so as to end up close to their notional equilibrium partitioning.

Partitioning enables validation to occur by reducing the large number of states that the unpartitioned historical data would demand, but partitioning does not come without a cost: the loss of some information. Whether the lost information is important can only be answered by examination of lost opportunities for profit-making on the part of the simulated firms, since, first, information is not an end in itself, but only a means to performing better in the oligopoly markets such as those we have discussed here, and, second, despite the attractiveness of entropy as a measure of information, there can be no one-dimensional measure of informativeness.

None the less, we have discussed three measures that might be used in a first cut at examining almost optimal partitioning: the number of perceived states, a player-specific measure, and, of course, entropy. Ideally, we should consider the impact on the firm's profitability of changes in the partitioning scheme. But this must await further research.

8. References

[1] Bertsekas D. P. & Tsitsiklis J. N. 1996, *Neuro-Dynamic Programming*, Belmont, Mass.: Athena Scientific.

[2] Blackwell D. & Girshick M.A. 1954, *Theory of Games and Statistical Decisions*, New York: John Wiley.

[3] Kalai E. & Stanford W. 1988, Finite rationality and interpersonal complexity in repeated games, *Econometrica*, **65**: 397–410.

[4] Lave C.P. & March J. G. 1975, *An Introduction to Models in the Social Sciences*, N.Y.: Harper & Row.

[5] Lipman B.L. 1995, Information processing and bounded rationality: a survey, *Canadian Journal of Economics*, **28**(1): 42–67.

[6] Marks R.E. 1992, Repeated games and finite automata, in *Recent Developments in Game Theory*, ed. by J. Creedy, J. Eichberger, & J. Borland, London: Edward Elgar, pp.43–64.

[7] Marks R.E. 1998, Evolved perception and behaviour in oligopolies, *Journal of Economic Dynamics and Control*, forthcoming.

[8] Marks R.E., Midgley D.F., & Cooper L. 1995, Adaptive behavior in an oligopoly, in *Evolutionary Algorithms in Management Applications*, ed. by J. Biethahn & V. Nissen, Berlin: Springer-Verlag, pp.225–239.

[9] McGuire C.B. 1972, Comparisons of information structures, Ch. 5 of McGuire & Radner (1972), pp.101–130.

[10] McGuire C.B. & Radner R. 1972, *Decision and Organization: A Volume in Honor of Jacob Marschak*, Amsterdam: North-Holland, 1972.

[11] Midgley D.F., Marks R.E., & Cooper L.G. 1997, Breeding competitive strategies, *Management Science*, **43**(3): 257–275.

[12] Mitchell M. 1996, *An Introduction to Genetic Algorithms*, Cambridge: MIT Press.

[13] Radner R. 1972, Normative theory of individual decision: an introduction, Ch. 1 of McGuire & Radner (1972), pp. 1–18.

[14] Radner R. 1987, Teams, in *Allocation, Information, and Markets*, ed. by J. Eatwell, M. Milgate, & P. Newman, New York: Norton, pp. 295–304, 1989; first published in *The New Palgrave*, New York: Norton.

[15] Rubinstein A. 1998, *Modeling Bounded Rationality*, Camb., Mass.: MIT Press.

[16] Simon H.A. 1972, Theories of bounded rationality, Ch. 8 of McGuire & Radner (1972), pp.161–188.

[17] Slade M.E. 1992, Vancouver's gasoline-price wars: an empirical exercise in uncovering supergame strategies, *Review of Economic Studies*, **59**: 257–274.

[18] Theil H. 1981, The maximum entropy distribution: second progress report, and Appendix, Center for Mathematical Studies in Business and Economics, Reports 8119 and 8120, University of Chicago, Department of Economics and Graduate School of Business.