

Evolved Perception and Behaviour in Oligopolies

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

To appear in
*The Journal of Economic
Dynamics and Control*

Presented at
Computing in Economics and Finance,
Geneva, 26–28 June, 1996.

Revised: May 23, 1997.

Please do not quote without the author's agreement. 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:

This paper builds on earlier studies which examined oligopolists in a repeated interaction as responding simply to past prices of their strategic rivals, and which used data from a mature market, with stable rules of thumb (mappings from past actions, or states of the market, to present prices) for the oligopolists' behaviour, whether purposefully learnt or emerging from the natural selection of the rivalry. The earlier studies imposed exogenous partitions on the action space, as perceived by the players. This study explores how such perceptions might be endogenised. A firm answer to the question of how oligopolists partition 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 the oligopolists have chosen to be. We use data from a retail coffee market to examine the evolved optimal partitioning and mapping of price space, manifest as the oligopolists' rules of thumb. The data suggest that brand managers are using very little information: whether prices changed or not.

ACKNOWLEDGEMENTS:

The author acknowledges many stimulating discussions with David Midgley, and his suggestion that Theil might have something to say about information and partitions; he also acknowledges Lee Cooper's provision of the data used. An anonymous referee made helpful suggestions. Research support was provided by the Australian Graduate School of Management and the Australian Research Council. An earlier version was presented at the "Commerce, Complexity, and Evolution" conference, University of New South Wales, February, 1996.

KEYWORDS:

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

1. Partitioning of Prices and History

This study follows from earlier work of 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 Midgley & Cooper 1995; Midgley Marks & Cooper 1997). That study used a genetic algorithm to derive “good” mappings from past actions to the marketing actions (such as price) to be pursued in the present period on the part of three strategic rivals. In the course of the earlier study 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. This study explores how such perceptions might be endogenised. A firm answer to the question of how players partition¹ 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.

In the U.S. retail ground coffee market the price has been seen to vary from about \$1.50 per pound to about \$3 per pound. Cluster analysis shows that some prices and price regions are used more frequently than others for each brand, and so the earlier study used four of these for each brand as the four actions in its simulation. Even though our earlier studies constrained the artificially intelligent adaptive economic agents to these four actions (of the many they had been seen to use, and of the 150 or more that are feasible), we found that the historical profit performances could be improved using simple four-action, one-round-memory machines.

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.²

We first formalise the process that each player uses in deciding his action in the market from one week to the next, using a framework outlined

-
1. 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), or earlier (Bohnenblust Shapley & Sherman 1949). Radner (1972) describes a model in which there is no strategic uncertainty, and in which our partitions are his signals.
 2. There is a literature on partitions as a means of imparting information to Bayesian-rational agents (Geanakoplos 1989, Samet 1990). A recent paper (Dimitri 1993) considers sequential experiments, in order to model certain information-processing skills on the part of agents.

by Lipman (1995). Each week, faced with the actual external state (or E-state), the players perceive an internal state (or P-state), which may update their beliefs, on which are conditioned their actions for the week, which together with the actions of their strategic rivals determine their profits.

2. Formalities

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 charged 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.

But it is unlikely that players perceive the information partition as finely as it is 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 $\zeta: \Omega \rightarrow \Theta$, that tells which *perceived* P-state 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 $\zeta(\omega)$.

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 or she 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, although for the moment we shall use a distinct set. 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): $u: A \times \Omega \rightarrow \mathbb{R}$, which describes how the state affects the value of the different actions available to any player. Note that here Ω represents the true E-state of the market during the present week, and will not be known 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 our case, discussed below, 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 $\beta: \Theta \rightarrow \Delta$ is the belief function. Beliefs matter because actions are contingent on them. The mapping from belief to action $\alpha: \Delta \rightarrow A$ is the action function.

Figure 1 illustrates the model from external E-states to final payoff, with an example of an E-state, showing how it is transformed into a P-state and an action. As Lipman points out, it is possible to combine perception,

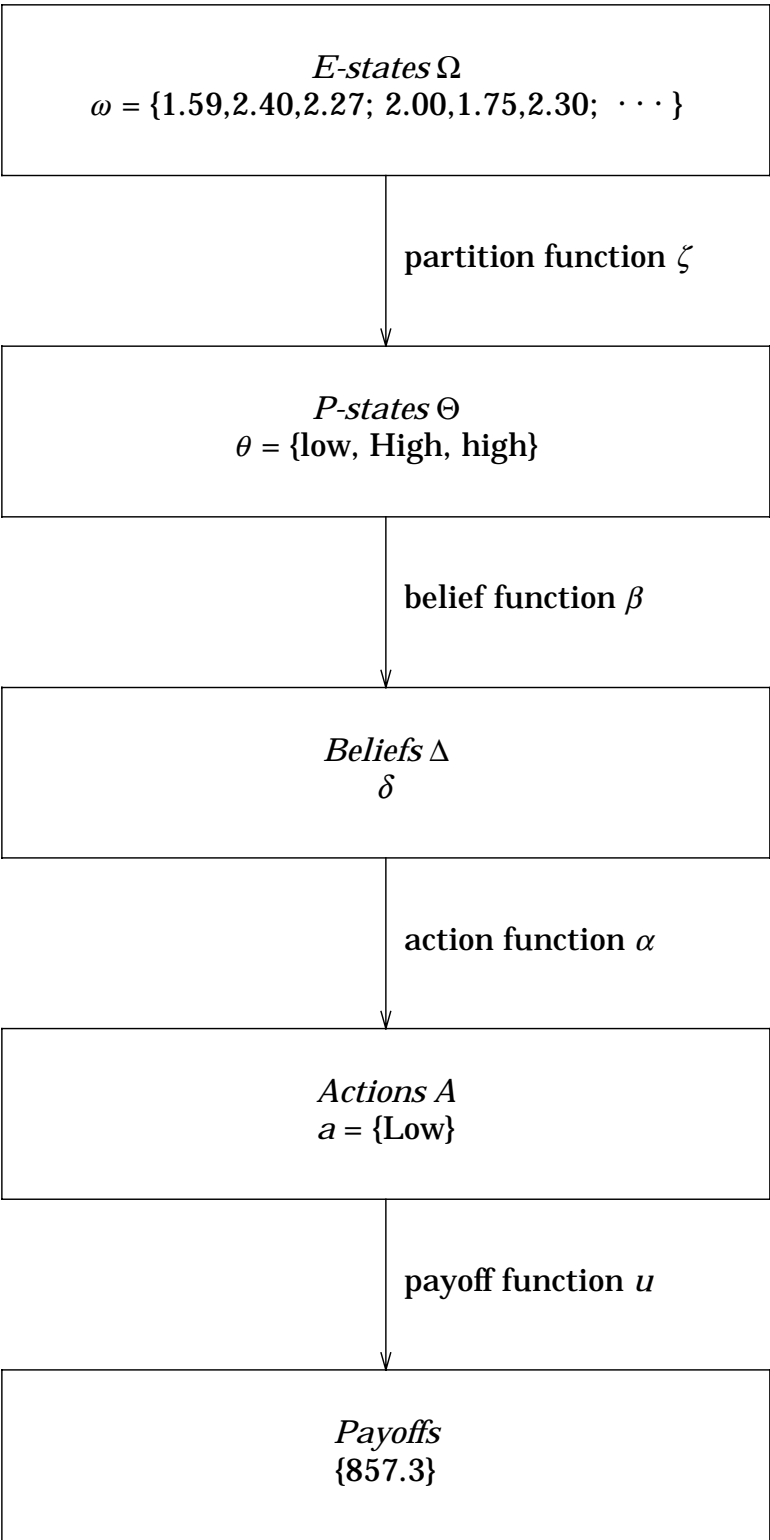


Figure 1: From External State to Payoff: The Player Modelled

belief, and action functions into a behaviour rule: $f: \Omega \rightarrow A$, where $f(\omega) = \alpha(\beta(\zeta(\omega)))$. This corresponds in the figure to an arc from the E-states node to the Actions node. In our earlier studies (Marks et al. 1995; Midgley et al. 1997), the partitioning mapping from the E-states node to the P-states node was exogenously determined, and we used a Genetic Algorithm (Mitchell 1996) to search for arcs (mappings) from the P-states node to the Actions node, using the payoffs as fitnesses.

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 $\alpha: \alpha(\delta)$ which maximises the following function for all δ must be derived:

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

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 evolutionary fitness 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 .

The belief function $\beta: \Theta \rightarrow \Delta$ can also be endogenous: if f is an interim-optimal behaviour rule $f: \Omega \rightarrow A$, then let V be the ex-ante expected profit associated with processing information according to the function β :

$$V(\beta) = \sum_{\omega \in \Omega} u(f(\omega), \omega) q(\omega).$$

Recall that β affects f by imposing a constraint on the set of available f s via the measurability requirement. Then we can model the determination of β by supposing that the player chooses it (endogenously) so as to maximise $V(\beta) - c(\beta)$, where c gives the expected information processing costs. Typically, the player is constrained to choose β from a set B of belief functions.

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. 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.

Partitions are closely related to equivalence relations: binary relationships that are reflexive, symmetric, and transitive. Given any equivalence relation R (such as two E-states ω and ω' are equivalent if they lead to the same beliefs), we can partition Ω into equivalence classes: for any given $\omega \in \Omega$, let

$$R(\omega) = (\omega' \in \Omega \mid \omega R \omega').$$

Thus $R(\omega)$ is the set of points equivalent under R to ω — the equivalence classes of R — which it is readily shown form a partition, Ω , called the partition induced by R . Similarly, given any partition Π , we can define an equivalence relation induced by the partition by saying that ω and ω' are equivalent if they are contained in the same event in Π .

Consider the equivalence relation over Ω defined by saying that two E-states ω and ω' are equivalent if they lead to the same beliefs (or if $\beta(\zeta(\omega)) = \beta(\zeta(\omega'))$). Let Π denote the partition induced by this equivalence relation, and let $\pi(\omega)$ denote the event of Π containing ω .

The key to the partitional models is that beliefs are assumed to be internally consistent, so that if the player's beliefs are δ only in certain external E-states, then δ must rule out any other E-states. Usually, the stronger assumption is made: that δ is assumed to be calculated from the prior q via Bayes' Rule. The partition can be used to summarise the player's information processing without explicit reference to the underlying belief function.

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.

4. Finite automata

Aumann (1981), Neyman (1985), and Rubinstein (1986) were the first to study repeated games in which players were restricted to using finite automata to implement their strategies. 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). Finite machines must by definition be bounded, and can be used to model the concept. Neyman and Rubinstein modelled bounded rationality as limitations on the number of states of the machine: Neyman imposed an exogenous limit on the number of states; Rubinstein assumed a cost trade-off.

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, $\sigma | h$. We can say that two histories, h and h' , are equivalent under σ if they lead to the same induced strategy; $\sigma | h = \sigma | h'$.

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 (Moore 1956). Banks & Sundaram (1990) considered the number of states and the number of transitions, in a two-dimensional measure.

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, following Fader & Hauser’s 1988 study. 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 a statistical technique on historical data of the oligopoly, namely cluster analysis, in order to partition the price space into four bands, again an arbitrarily chosen number. The boundary prices 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 the machine-learning genetic algorithm 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 can define our response function as a mapping from P-state to action: $\gamma: \Theta \rightarrow A$. See our earlier studies (Marks et al. 1995; Midgley et al. 1997) for discussion of this search.

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 $f(h) = \sigma | h$. 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.³ 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.

A related model is Dow's 1991 model of search with limited memory. The E-state is a pair of prices, one observed in period 1 and one in period 2. The action is which price at which to buy. Dow's agent knows (or believes) he will not be able to remember the first price exactly (an exogenous constraint) and so is modelled as partitioning the set of possible prices; his memory is only into which event of the partition the period 1 price fell. The partition is explicit, and Dow also assumes that costs are proportional to the number of events in the partition.

The motivation for this study is the desire to endogenise the partitioning which is necessary for simulations of interactions as mappings from state of the market, suitably defined, and actions by the players. Although this study does not go beyond the discussion below of how to partition to maintain maximal information, future work will harness these results in further simulations, as discussed above.

5. Optimal Partitioning

Lipman 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 a finer partition. Although the

3. Lipman (1995, fn. 5) points out that not every partition of I can be generated by some strategy, and that not every function from such a partition to A will constitute a legitimate strategy.

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.

programming effort increases in the number of thresholds, Lipman reports that this class of models assumes zero cost for information processing.

5.1 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 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. The data are 78 weeks of weekly observations for three competing brands in a mature market (that for canned, ground coffee in a U.S. city).⁵

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.

We also consider dichotomous partitioning of the first differences of price (both absolute and algebraic) in order to model players’ responses to price jumps, as well as considering a symmetric terchotomous partition of price levels.

5.2 Measures of Optimality

Which partitioning is best?

We argue that the best partition is the one that loses the least amount of information. 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. For further details, see Midgley Marks & Cooper (1997). The three rivalrous brands are Folgers, Maxwell House, and Chock Full O Nuts.

6. As discussed in Section 7 below, 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 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”. Note that Kolmogorov introduced the concept of the entropy of a countable partition in 1958 (Iyanaga & Kawada 1977).

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

where there are N perceived states, and the probability (or observed frequency) of state i is p_i . If the logarithm base b is 2, then the units of entropy are “bits”; if b is the exponential constant e , then “nits”.

Theil argues on axiomatic grounds that entropy is justified, by showing that the entropy information measure of an event (such as the observation of a specific market state) satisfies four axioms:

Axiom 1. The information content of observing that a state i occurred depends only on the probability p_i of its occurrence prior to the observation.

Axiom 2. The information is a continuous function of p_i in $0 < p \leq 1$ and monotonically decreasing.

Axiom 3. When state i is certain, its observation carries zero information.

Axiom 4. The information content of a state which is the union of two mutually exclusive states (zero intersection) equals the sum of the information of observing one state and that of observing the other. (Additivity.)

The maximum number of 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:

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

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”.

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.

Figure 2 uses the market data from Chain One⁷ to plot two measures of information losses against the position of a dichotomous partition, as it moves in steps of a thousandth of the range between lowest price and highest price charged by each of the three brands over the 78-week period of the data. The

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.

The steps we follow are:

1. For each of the three brands or players, determine the minimum and maximum prices charged over the period.
2. For the following five steps, choose a partition point λ , ($0 \leq \lambda \leq 1$). With iteration of these steps, the partition point will increment from zero to one in steps of a thousandth of the range between each brand's minimum and maximum prices.⁸
3. For each brand's price, for each week, for a given partition point λ , determine whether the price is "low" or "high": if
price \leq minimum price + λ (maximum price – minimum price),
then the price falls "low", otherwise "high". When $\lambda = 0$, almost all prices are classified as "high", since the partition point is at the minimum price, and so the "low" set is almost empty; when $\lambda = 1$, all are classified as "low", for the opposite reason. (Note that the minimum and maximum prices in the expression are brand- or player-specific, since there is no constraint on players to price in the same range as do their rivals.)
4. For the given partition point, determine the perceived state of the market. With three players and one-week memory, there are 2^3 or 8 possible states. Arbitrarily number them by calculating the state number:

$$\text{state} = 4 \times F + 2 \times MH + CFON,$$

where F is Folgers' action, MH Maxwell House's, and $CFON$ Chock Full O Nuts'; if a player priced "low" last week, then define that player's action to be 0, otherwise 1. (For Figure 4, with two-week memory, there are 2^6 or 64 possible states, similarly arbitrarily numbered.)

5. For the given partition point, calculate the observed frequency p_i of each state i , as defined by the number of times each perceived state is observed as a proportion of the total number of times states are perceived.

7. We had scanner data from three supermarket chains, but here use data from Chain One and (in Figure 3) Chain Two only.

8. With the data we have, increments smaller than one thousandth do not reveal any finer structures in the entropy measure: one thousandth is a sufficiently fine increment.

6. For the given partition point, these frequencies can be thought of the sample probabilities of the perceived states, which enables calculation of the *alog entropy measure* associated with a particular partition point, from equation (3).
7. As mentioned above, the the total number of perceived states would be equal to the alog entropy if the frequencies of all perceived states were equal: using *the total number of perceived states* as a measure of the effectiveness of a partition point at retaining *relevant* information means throwing away the information of the frequencies of the perceived states, for the given partition point. Alog entropy is bounded above by the the total number of perceived states, as seen for instance in Figure 2.
8. After incrementing the partition point by one thousandth, Steps 2 to 7 are repeated, until the full interval between minimum and maximum prices has been searched, as plotted in the Figures below.

We examine the possibility of price changes, as well as price levels, below, calculating both measures.⁹

6. Results

Following the eight-step procedure listed above, we use 78 weeks of data from Chain One to plot the measures in Figure 2. This is compared with Chain Two data in Figure 3. For the Chain One data, we have also considered two-week memory (Figure 4); first-differences in prices with one-week memory, both absolute, in which only the size of the price jump matters (Figure 5), and algebraic, in which both size and direction of the price jump matter (Figure 6); and a symmetrical terchotomous partition, in which the price range of each brand is divided into three partitions, symmetrically about the mid-point of the range, which requires a single parameter only (Figure 7).

The six cases are summarised in Table 1 and the results are summarised in Table 2, below.

9. An earlier version of this paper explored the possibility of a third, brand-dependent measure. This was based on the concept of the mappings from state to action of a specific brand manager. It is possible to derive a matrix, the rows of which correspond to perceived states, given the partition point, and the columns of which correspond to a fixed number of price ranges which span the actions of a specific player. Intuitively, when the partition is such as to maximise the number of perceived states, the best partition is that which minimises the mean mapping from perceived state to action. The rationale is that, under the ideal partition, all possible states are perceived, and that each state is found to map to only one action in the historical data. If, in the limit, it is found impossible to reduce the number of actions per perceived state to one, with all states perceived, then this may be due to one of several possibilities: a misspecification of the model (it may be, for instance, that players respond to not price levels but to price changes), or that the assumption of a deterministic mapping from state to action is wrong, with some mixing of strategies. We do not explore this further here.

Figure	Data (Chain)	Partitions	Weeks of Memory	Price Variable
2	One	2	1	level
3	Two	2	1	level
4	One	2	2	level
5	One	2	1	difference
6	One	2	1	difference
7	One	3	1	level

Table 1: Summary of the Cases.

6.1 Dichotomous Partition of Price Levels

6.1.1 One-week memory. Examination of Figure 2 reveals that entropy is a much finer measure than is the number of perceived states: over ranges of the partition point λ , the latter measure is unchanging, while the entropy varies. Moreover, from Figure 2 the entropy measure suggests that the optimal dichotomous partition for Chain One data is a threshold λ at 90.1% This means that for a dichotomous partition with the observed behaviour of the three strategic brands in Chain One, less information is lost with a partition point at 90.1% of the distance between each brand's minimum and maximum prices than with any other partition point.

For Chain Two data, the optimal partition point using the entropy measure is significantly lower, at between 71.4% and 71.5%, as seen in Figure 3. Comparing Figures 2 and 3 for the two sets of data, we can characterise the players in Chain Two as more responsive to prices in the mid-upper range than the players in Chain One appear to be. In Chain One, the threshold between "high" and "low" prices is around 90% of the price range between highest and lowest price of each player; in Chain Two it is around 71%.

6.1.2 Two-week memory. Players with two-week memories can respond to *movements* in other players' prices, unlike players with one-week memories only: a rise (from "low" to "high"), a fall, a steady "low", or a steady "high" on the part of each of the other players (as well oneself) can be responded to. With two-week memory, three players, and dichotomous partitioning, there are $2^6 = 64$ possible states of the market. Comparison of Figures 2 and 4 reveals a richer information structure of the same data set (Chain One). But the entropy-maximising partition point λ is exactly the same: at 90.1% precisely (using steps of 0.001 of the range). This is only an artefact of the data, as is seen when the Chain Two data are analysed with a two-week memory model: the maximum entropy partition is slightly higher than in Figure 3, at 72.4%.

6.2 Dichotomous Partition of First Differences of Price

Using two-week memory the model can track the directions of players' price movements, but not, with dichotomous partitioning, the magnitudes of price




Figure 2: Information Measures: Perceived States and Entropy




Figure 3: Information Measures: Perceived States and Entropy




Figure 4: Information Measures: Perceived States and Entropy

jumps. For this reason, and since it may better describe the way the historical brand managers behaved, we consider not price levels, but the *first differences of prices*: price at week t minus price a week earlier.

Figure 5 deals with absolute first differences, so that only the magnitude but not the direction of the price jump matters, and Figure 6 deals with the actual (or algebraic) first difference, which include both magnitude and sign of the change. For Figure 5, the partition point is between 0% and 100% of the range between the smallest (greater than or equal to zero) and largest absolute first difference in price for each player; for Figure 6, between 0% and 100% of the range between the smallest (most negative) and largest (most positive) first difference for each player.

6.2.1 Absolute First Differences. From Figure 5, we see that entropy is maximised at $\lambda = 1.4\%$ and 1.5% (with steps of 0.001), which means that with a dichotomous partition the most significant threshold is whether or not the absolute value of a price change is greater than or less than about 1.4% of the range between smallest absolute price change and largest. Intuitively, the threshold is the boundary between no change (or negligible change) in price from one week to the next, and a significant change (which is here any absolute change greater than 1.6%, up or down).

6.2.2 Algebraic First Differences. Figure 6 is plotted for actual first differences in price. As we might expect, it is more closely symmetrical than is Figure 5, since price rises and falls register distinctly here. Maximum entropy occurs at a λ equal to the exact midpoint of the range of price differences. This seems consistent with the results of Figure 5, but there is no reason, *ex ante*, to believe that price rises and falls should be even roughly symmetric: a pattern of small falls followed by large rises — such as is sometimes seen in petrol price wars (Slade 1992) — will bias the data, and the midpoint will correspond to a positive price jump, but such asymmetries are not seen in these data.

6.3 Symmetrical Terchotomous Partition of Price Levels

A dichotomous partition is the simplest we can consider (and the easiest to calculate: above or below the partition point); with only one degree of freedom, it is also the easiest to plot, as in Figures 2 through 6. But it may be that players use more sophisticated partitioning of the price space. For terchotomous (3) or higher-order partitions, there are more than one degrees of freedom, in general, which is more difficult to search for and not as easy to present graphically. There is, however, one way to model terchotomous partitioning using only one degree of freedom: λ is the proportion of the range spanned by the central partition of the three, centred on 50%, so that $\lambda = 0$ corresponds to two partition points together at the centre point of the price range; $\lambda = 1$ corresponds to one partition point at the bottom (left) of the price range and the other at the top (right); $\lambda = 50.0\%$ corresponds to one partition point at the quarter point and the other at the three-quarters point of the price range — hence the description *symmetrical terchotomy*.

Figure 7 presents the results of using this partitioning with the Chain One data. The maximum entropy threshold λ occurs at 80.3%. This means




Figure 5: Information Measures: Perceived States and Entropy




Figure 6: Information Measures: Perceived States and Entropy




Figure 7: Information Measures: Perceived States and Entropy

that when the price range is divided into three partitions at about 10% and 90%,¹⁰ then least information is lost. That is, the most informative symmetrical terchotomous partition is sensitive to price levels near the bottom (aggressive) or near the top, and doesn't differentiate over the large range of prices in the middle (80% of the range). Of course, there is no reason why the players should perceive price symmetrically in this manner.

6.4 Maximum Perceived States

In almost all the cases considered above, the partition which minimises the entropy measure of information loss occurs in or at the boundary of the region of maximum perceived states. This is summarised in Table 2, below, which characterises the two measures of antilog entropy and maximum perceived states in terms of the optimal partition point λ and the two respective measures, in units of effective states.

Figure	Entropy		Perceived States	
	λ (%)	measure (states)	λ (%)	measure (states)
2	90.1	6.525	86.5–93.8	8
3	71.4–71.5	6.235	54.3–71.5, 72.9–76.0	8
4	90.1	22.889	90.2–90.7	30
5	1.4–1.5	7.765	0.3–15.6	8
6	50.0	4.663	48.8–49.9	8
7	80.3	11.334	81.8–82.7	19

Table 2: Points of Maximum Information Retained

Figure 7 presents the greatest distance between the optimal partition points of the two measures: 80.3% for maximal entropy is 1.5 percentage points below the range of 81.8% to 82.7%. This reveals that, especially as the number of possible states rises, it is not necessary that the maximum-entropy point occur within the range of maximum perceived states.

In the partitioning models of Figures 2, 3, 5, and 6, the maximum number of states is 8, since there are three players, each with two possible prices, “high” or “low”. We see that there exist partition points which allow all 8 states to be perceived from the Chain One data (and, from Figure 2, the Chain Two data). Figure 4 is plotted from a model with a maximum number of 64 states, and with the Chain One data the maximum perceived is 30. Figure 7 is plotted from a model with a maximum number of 27 states, since each of the three players can have three possible prices; with the Chain One data the maximum perceived is 19.

10. 10% = 50% – 80% of 50%; 90% symmetrically.

6.5 Most Information Retained

The results of the six models, summarised in Table 2, provide sufficient information for a tentative answer to a question posed at the outset: how much information do the brand managers choose to use in their repeated interactions? Since we have not here developed a brand-dependent measure, we use instead the two measures of information retained — maximum perceived states, and entropy — with various models of simple partitioning. The individual information-processing of the managers has resulted in three, interrelated series of actions over 78 weeks, which we have analysed, and which can reveal which simple partitioning is best at retaining information.

As shown, the maximum possible number of perceived states provides an upper bound to the entropy. This allows us to compare three of the models, those underlying Figures 2, 5, and 6, all of which allow a maximum of eight possible perceived states. The other models are difficult to compare because they use different data (Figure 3), or have different numbers of possible states (Figure 4, up to 64, and Figure 7, up to 27).

All three of the Figures 2, 5, and 6 reveal the maximum number of eight perceived states. They can be ranked in terms of maximum entropy as Figure 5 (7.765/8), Figure 2 (6.525/8), and Figure 6 (4.663/8). Our tentative conclusion is that, using a dichotomous partition of the price actions, the brand managers are revealed as using a very simple model. The model underlying Figure 5 with a partition point of 1.4%, as discussed in Section 6.2.1, means that if the managers respond simply to whether their rivals changed their prices or not, comparing actions two weeks ago with actions one week ago (first differences), then they glean more information from the historical data than if they ask whether the prices were raised or lowered (Figure 6) or whether the prices last week were “high” or “low” (Figure 2).

The Figure 5 model implicitly uses two-week memory to derive the first differences of price; the Figure 4 model uses two-week memory explicitly. Its maximum entropy measure of 22.889 is not close to the maximum number of perceived states of 30 (or a maximum possible of 64), which leaves the model of absolute first differences of prices (of Figure 5) with the highest information retained under the dichotomous partition point of 1.4%. Similar analysis of the historical data from Chain Two (unplotted, but a partition point of 2.3% and an entropy of 7.803 effective states) confirms that the simple dichotomous model of price change or not shows the highest level of information retained.

7. Conclusion

This paper has reported the latest stage in our programme of using the techniques of machine learning and the historical market data at our disposal to compare the behaviour of our artificially intelligent adaptive economic agents using evolving partitions with the behaviour of such players using exogenously determined partitions. We have shown how our search space has been widened from the mapping of the response function $\gamma: \Theta \rightarrow A$ (from P-state to action) to the mapping from E-state to action: the behaviour rule $f: \Omega \rightarrow A$. We have discussed and demonstrated techniques for measuring the

information retained associated with a particular partition. Using the entropy measure we have examined optimal dichotomous partitions of historical data of price competition, and derived optimal partition points. This is helping us us to answer the question, at least for our historical data, of how boundedly rational our players, the coffee brand managers, were.

Our conclusion, based on the simple partitioning models considered, and the historical data of the three-way interactions between the brand managers, is that the simple model of whether or not prices changed is the most informative fit with the two sets of historical data examined. This implies that the boundedly rational historical brand managers used this simple model in their decisions to respond, whether consciously or not we cannot say.

We have considered optimal partitions using entropy as the appropriate measure of information retained. But the information is not an end in itself, as we have implicitly assumed. Rather, the information is a means to an end — to maximise net return in the repeated interaction among the players.¹¹ The next stage of the research program is to search for the best combination of information partitioning and the consequent mapping from state to action, using exogenous actions, as previously.¹² 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, action mappings, and final actions in the repeated interactions. But this must await future work.

8. References

- [1] Aumann R. 1981, Survey of repeated games, in: *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, by R.J. Aumann et al., Zurich: Bibliographisches Institut, pp.11–42.
- [2] Banks J.S. & Sundaram R.K. 1990, Repeated games, finite automata, and complexity, *Games and Economic Behavior*, **2**: 97–117.
- [3] Blackwell D. & Girshick M.A. 1954, *Theory of Games and Statistical Decisions*, New York: John Wiley.
- [4] Bohnenblust H.F., Shapley L.S., & Sherman S. 1949, Reconnaissance in game theory, Research Memorandum RM-208, RAND Corporation.
- [5] Dimitri N. 1993, Learning partitions, *Economics Letters*, **42**: 195–199.
- [6] Dow J. 1991, Search decisions with limited memory, *Review of Economic Studies*, **58**: 1–14.

11. 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.

12. Radner’s information function and decision function, respectively.

- [7] Fader P.S. & Hauser J.R. 1988, Implicit coalitions in a generalized Prisoner's Dilemma, *Journal of Conflict Resolution*, **32**: 553–582.
- [8] Geanakoplos J. 1989, Game theory without partitions and application to speculation and consensus, Discussion paper No. 914, Yale University.
- [9] Iyanaga S. & Kawada Y. (ed.) 1977, *Encyclopedic Dictionary of Mathematics*, by the Mathematical Society of Japan, Cambridge: M.I.T. Press, Section 146E, p.493.
- [10] Kalai E. & Stanford W. 1988, Finite rationality and interpersonal complexity in repeated games, *Econometrica*, **65**: 397–410.
- [11] Lipman B.L. 1995, Information processing and bounded rationality: a survey, *Canadian Journal of Economics*, **28**(1): 42–67.
- [12] 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.
- [13] 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.
- [14] Marschak J. & Radner R. 1972, *Economic Theory of Teams*, New Haven: Yale Univ. Press.
- [15] McGuire C.B. 1972, Comparisons of information structures, Ch. 5 of McGuire & Radner (1972), pp.101–130.
- [16] McGuire C.B. & Radner R. 1972, *Decision and Organization: A Volume in Honor of Jacob Marschak*, Amsterdam: North-Holland, 1972.
- [17] Midgley D.F., Marks R.E., & Cooper L.G. 1997, Breeding competitive strategies, *Management Science*, **43**(3): 257–275.
- [18] Mitchell M. 1996, *An Introduction to Genetic Algorithms*, Cambridge: MIT Press.
- [19] Moore E.F. 1956, Gedanken-experiments on sequential machines, in: *Automata Studies*, ed. by C.E. Shannon & J. McCarthy, Princeton: Princeton Univ. Press, pp. 129–153.
- [20] Neyman A. 1985, Bounded complexity justifies coöperation in the finitely repeated Prisoners' Dilemma, *Economics Letters*, **19**: 227–229.
- [21] Radner R. 1972, Normative theory of individual decision: an introduction, Ch. 1 of McGuire & Radner (1972), pp. 1–18.
- [22] 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.
- [23] Rubinstein A. 1986, Finite automata play the repeated Prisoner's Dilemma, *Journal of Economic Theory*, **39**: 83–96.

- [24] Samet D. 1990, Ignoring ignorance and agreeing to disagree, *Journal of Economic Theory*, **52**: 190–207.
- [25] Simon H.A. 1972, Theories of bounded rationality, Ch. 8 of McGuire & Radner (1972), pp.161–188.
- [26] Slade M.E. 1992, Vancouver's gasoline-price wars: an empirical exercise in uncovering supergame strategies, *Review of Economic Studies*, **59**: 257–274.
- [27] 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.