Journal of Behavioral Decision Making
J. Behav. Dec. Making, **13**: 19–34 (2000)

Do Consumers Ever Learn? Analysis of Segment Behavior in Experimental Markets

JAIDEEP PRABHU[1]* and GERARD J. TELLIS[2]
[1] *Tilburg University, The Netherlands*
[2] *Marshall School of Business, University of Southern California, USA*

ABSTRACT

Economists argue that, despite cognitive limitations, economic agents arrive at optimal choice rules by learning. The assumption is that consumers, for example, are adaptively rational. Adaptive rationality raises a host of issues. We address three of these in the context of experimental markets: do consumers differ on the basis of learning; how do these differences, when aggregated, affect market efficiency; and how do consumers learn? Analysis of our experimental data reveals the following. First, multiple segments of consumers exist on the basis of learning. Second, the largest segment consists of subjects who do not learn despite timely feedback and motivation. Third, although some consumers do learn to make optimal choices, the effect of this segment on market efficiency is cancelled by an equal number of subjects who 'learn' false relations. Finally, although subjects do not learn strict rationality even with experience, they are in the aggregate not so irrational as to allow highly suboptimal brands to survive. Further analysis of how consumers learn, specifically on the cues (signals) and the rules consumers employ in making choices over time leads to the following two conclusions. First, some signals make learning more easy than others: for example, providing market share information improves learning but not as much as providing quality information does. Second, people employ different rules depending upon the type of information they have. For example, consumers making decisions based only on price information are more likely to use a heuristic like 'buy a medium-priced product provided it has not failed in the past'. Consumers making decisions based on price and quality information may employ a heuristic such as 'buy top quality products regardless of price'. We discuss the implications of these findings for theory and practice. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS rationality; learning; market efficiency

INTRODUCTION

Economic theory asserts that agents in markets live by the dictates of strict rationality. When confronted with evidence that consumers are subject to cognitive errors and systematic biases, the theory

* Correspondence to: Jaideep Prabhu, Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands.

CCC 0894–3257/2000/010019–16$17.50

counters these findings with several arguments. One argument is that economic agents arrive at optimal rules, despite their limitations, by a process of learning, or by *adaptive rationality* (Alchian, 1950; Lucas, 1987). Besides being a plausible explanation of rational behavior, learning also provides a dynamic interpretation to markets that are in a state of evolution (Haines, 1969; Miller and Plott, 1985). The notion of adaptive rationality or dynamic behavior raises a host of questions: do consumers learn fully with experience? If yes, do they differ in their learning? If so, how do these differences, when aggregated over a market of consumers, affect market efficiency?

Behavioral researchers focus on different aspects of consumer learning to arrive at different conclusions from those of researchers in economics. For instance, while behavioral decision theory focuses on learning at the individual level, economic theory focuses on learning or its outcome at the market level. Casual analysis suggests that learning operates at two levels — the individual and the aggregate. Namely, while individuals learn, markets may do so too, i.e. they equilibrate and become more efficient over time. Learning at the individual level, at the aggregate level, and the interaction of the two, are important to all who study markets. We suggest, therefore, that analysis of segments of consumers with homogenous patterns of learning is an intermediate step that links the above two research traditions while also providing fresh insights.

The economic and behavioral traditions also differ in their broad conclusions. Behavioral decision theorists argue that individuals are subject to cognitive illusions that preclude rational behavior (see Tversky and Kahneman, 1986). They assert that individuals are susceptible to biases in choices that persist even in the presence of information and motivation. Economists, on the other hand, argue that markets may become rational in several ways, even though individual agents are not. In particular, Camerer (1992) discusses four of these mechanisms which he names the *cancellation*, *the smart few*, *the learning* and the *evolutionary mechanism*. Each of these mechanisms has a possible counter-mechanism, however, and the validity of each is therefore an empirical issue.

The goal of this paper is to study whether consumers learn, and if so, how that affects market efficiency. Towards this goal, we address three main issues. First, we focus on differences in learning by segments of individuals. By introducing segments as an intermediate level of analysis, we link the individual-level approach of behavioral decision theory with the market-level approach of economics. A segmentation approach has several other advantages. Market level analysis is subject to the problem of heterogeneity and the aggregation bias that results from pooling different types of individuals. Individual-level analysis, on the other hand, provides too much information to be easily grasped and interpreted. Segmentation gets around both problems. Indeed, analyzing segments of informed and uninformed consumers has been an important means for relaxing traditional assumptions of economic theory (e.g. Salop and Stiglitz, 1977). Further, marketers and public policy makers routinely invoke segmentation to account for a variety of market phenomena as well as to develop marketing strategies and policies.

Second, we study the impact of differences in segment learning on market efficiency. Specifically we test the validity of the mechanisms that Camerer suggests for market efficiency when individuals are not rational. In particular, we test the cancellation, the smart few and the learning mechanism. We do not test the evolutionary mechanism because that mechanism applies to sellers more than buyers. We do, however, test a combined outcome of the smart few and evolutionary mechanisms.

Third, we study the dynamics of consumer learning. In particular we study how information on quality and market behavior affects learning, how cues and rules affect consumers' choices over time.

We first develop a theory of learning at the individual, segment and market levels. Next we describe the details of a study conducted to test hypotheses based on the theory. We discuss the results of the study. Finally we present conclusions and implications for future research.

THEORY

We define quality as those attributes of which all consumers unambiguously prefer more. Examples include product reliability, timely delivery or pleasant service. In contrast, stylistic attributes are those over which consumers may differ in their preferences. Examples are color, design, or sweetness. We define the rational choice as that which maximizes the consumers's expected utility. In the context of products and services, this choice requires that consumers integrate information on quality and price of brands within categories to arrive at estimates of expected value. We define learning, therefore, as the process of identifying quality and integrating it with price information to make rational choices. We refer to the term market, here, primarily as the aggregate of consumers making choices. We define market rationality or efficiency as the degree to which the market chooses the best outcome. With these definitions, we discuss the theory of consumer learning at three levels: the individual, segment and market level.

Consumer learning at the individual level

Consumers may come to know product quality from personal experience or published literature or retail outlets. Alternatively, because consumers find it hard to observe the true quality of products, they may use cues or signals to infer quality (Spence, 1974; Spiglitz, 1975). They may use market share (Smallwood and Conlisk, 1979), word-of-mouth (Kennedy, 1994), advertising (Hoch and Ha, 1986; Milgrom and Roberts, 1986) or price itself (Monroe and Petroshius, 1981; Wolinsky, 1983) to infer product quality.

The problem with these cues is that each is an imperfect indicator of quality. Personal experience is limited by constraints of time and money: consumers cannot sample exhaustively or systematically in order to learn the best brand in an objective sense. Even if they could, because of limited recall and heightened sensitivity to recent experiences (Tellis and Gaeth, 1990), consumers may not process experience rationally. Search of the literature or retail outlets is similarly limited. Indeed, even though sources such as Consumer Reports may be accurate, few consumers refer to these sources or interpret them correctly. Word-of-mouth is likely to be only as good as the information of the individual consumers who generate it. Advertising information is biased because advertisers differentially exaggerate the quality of their products depending on their costs of advertising and quality (Tellis and Fornell, 1988).

The process by which consumers combine quality and price may be regarded as a problem in decision making under certainty. The rational choice for consumers is to maximize expected utility. In particular, when quality is uncertain, they need to combine the underlying probability distribution of quality with the costs of various levels of quality. A growing body of research indicates that people do not understand this rule of expected utility and do not learn it even when it is explained to them (Kahneman and Tversky, 1979). Thus, consumers may not combine price and quality optimally. Further, given that identifying quality and combining it with price are cognitively difficult tasks, and that errors in decision making may not always be perceptibly costly, consumers may consider a trade-off between effort and accuracy in making choices. Bettman, Johnson and Payne (1991) develop such an effort-accuracy framework and use it to discuss several rules or heuristics that consumers may use in making choices. For instance, consumers may use a lexicographic rule, i.e. choose the brand that has the best value on the most important attribute. An example of such a rule would be 'choose the highest price brand' or 'choose the brand with the highest quality'. Or consumers might employ a satisficing rule, i.e. consider brands one at a time and choose the first one that meets a minimum cut-off level on an important attribute. An example of such a rule would be 'choose a brand with quality that's above average'. While these rules are likely to reduce the complexity of the processing task, as with signals of quality they are likely to be error-prone as well, and may lead consumers to make irrational choices at least in the short run.

Consumer learning at the level of segments

Segments of both rational and nonrational consumers may concurrently exist in markets. Kahneman and Tversky (1979) and Kahneman, Knetsch and Thaler (1987) each reported data on choices designed to test for rationality, defined as behavior that maximizes subjective expected utility. These authors emphasized that in most cases a majority of subjects did not make the rational choice. They glossed over the fact, however, that in some cases between a quarter and a third of the subjects disagreed with the majority and did make a choice that was rational. How stable were these rational and 'irrational' behaviors? Were the same people being rational all or most of the time? Who were these people? The authors did not pursue whether such rational behavior was systematically linked to certain individuals or segments.

We hypothesize that at least four types of consumer segments exist (see Exhibit 1). First, given the latent nature of produce quality, the complexity of markets, and the unsystematic manner in which consumers obtain and integrate information, some consumers may never learn to make optimal choices (nonlearners). Second, some consumers may not make good decisions when they first start, but through trial and error, they may learn optimal rules through experience (learners). Third, some consumers may both gather the right information and make the right decisions from the start (early learners). Finally, some consumers may 'learn' incorrect or suboptimal rules of choice, i.e. their decisions may systematically deteriorate over time (confused learners). The relative sizes of these segments as well as their interaction with each other will determine the extent to which the market as a whole learns or approaches efficiency.

The impact of consumer learning on market behavior

Markets are an aggregation of individual consumers. In a static sense, the proportion of rational consumers at any time determines the extent to which the market is efficient at that time. In a dynamic sense, the proportion of consumers who learn to make the optimal choice and the time they take to do so determines the extent to which the market will be efficient and reach equilibrium.

Camerer (1992) discusses four mechanisms by which markets may be rational or efficient, even though individual agents are not. His discussion is general and probably applies to producers more

Exhibit 1. Segments of consumers based on patterns of learning

Segment type	Pattern of learning behavior
Nonlearners	Consumers who make suboptimal decisions throughout their purchase history and do not register a significant movement towards or away from optimality with time
Learners	Consumers who begin with suboptimal choices but move towards optimality over time
Early learners	Consumers who make optimal choices from the start of their purchase histories and who continue to do so over time.
Confused learner	Consumers who 'learn' false relationships or consistently apply spurious cues and decision rules and thus move away from optimality over time

Exhibit 2. Mechanisms and counter-mechanisms for market efficiency despite individual irrationality

Mechanism	Logic of mechanism	Logic of counter-mechanism
I Cancellation	Errors cancel because they are randomly distributed about the efficient choice	Errors are systematic and will not cancel in the aggregate
II Smart few	The activity of a few rational consumers wipes out the errors of less rational ones	What if experts also make systematic errors?
III Learning	Less rational consumers learn from experience or from rational others	Learning is difficult in complex environments; information transfer across people and contexts is inefficient
IV Evolution	Less rational consumers are selected out	Stakes are not high enough to eliminate irrational consumers; new irrational consumers are always entering markets
V Hybrid	Highly irrational consumers are few; market share of highly suboptimal brands is insufficient for them to remain in the market	Irrational consumers may be numerous enough to support even the most suboptimal brands

Adapted from Camerer (1992).

than consumers. This paper, however, applies his categories to the decisions of consumers. We discuss these mechanisms below. They are also summarized in Exhibit 2.

- *The cancellation mechanism.* Logic: If individual choices are randomly distributed about the rational choice, then the errors of some will be cancelled by the good choices of others and the market as a whole will appear efficient. Counterargument: Individual errors are often systematic. Further, because markets are aggregations of individual agents, individual errors when summed will compound the inefficiency of the market.
- *The smart few mechanism.* Logic: A few smart consumers will adjust their behavior to take advantage of the stupidity of others, resulting in an efficient market. Counterargument: Smart consumers may not wish to take advantage of irrational others. Even if they are inclined to do so, the smart few may not be able to exploit irrational consumers in a free consumer market. What is more likely is that irrational consumers will follow the behavior of rational ones and learn vicariously through them. Stated thus, the smart few mechanism seems a plausible one, although it is possible that even the smart few are prone to systematic errors or that such consumers are not easily identifiable in the marketplace. In which case, the smart few mechanism is likely to have a weak effect on improving market efficiency.
- *The learning mechanism.* Logic: Consumers can learn to avoid errors with experience or from knowledge gained from other more rational consumers. Counterargument: Learning from experience is difficult for a number of reasons as described above (see section on learning at the individual level; also see Einhorn and Hogarth, 1978). Further, less experienced consumers may rely on more experienced ones for help or on cues such as market share, but these are by no means foolproof: information exchange across contexts and people may itself be inefficient.
- *The evolutionary mechanism.* Logic: Traders who are persistently irrational go bankrupt and are eliminated leaving only rational traders in the market. This argument, in Camerer's original formulation, is not directly relevant to the consumer context. A hybrid of the evolutionary mechanism and the smart few mechanism is however relevant to this study and is considered below.
- *A Hybrid Mechanism.* Logic: Suppose the number of consumers who make highly suboptimal decisions is so small that highly suboptimal brands achieve only a very small market share. Further, suppose that a firm must achieve a minimum market share to remain viable. In this case, highly suboptimal brands will be eliminated from the market even though some highly irrational consumers

(the stupid few) exist. This argument is a variant of the argument presented by Salop and Stiglitz (1977). Counterargument: The number of highly irrational consumers may be large enough to keep suboptimal brands viable over a significant period of time.

We next proceed to test these alternate mechanisms through experimental simulation of consumer choice in contemporary markets.

METHOD

Design

We developed a computer-based market game to provide a dynamic context in which subjects could make product decisions over time. The game lasted 24 periods. In each period subjects chose to fly from LA to London and back by one of ten possible airlines. All subjects had the names of the airlines,[1] the price of tickets on each airline, as well as the cost they (subjects) would incur if their flight was delayed. We set cost of a delay at $700, regardless of the airline or the individual. We explicitly informed consumers not to consider other attributes of airlines such as service and frequent-flyer miles. Thus we let on-time performance be the only measure of quality. After each period, subjects learnt if the chosen airline was on time or not. If the flight was delayed, the cost of the delay together with the price of the ticket was added to their total costs for that period. Subjects then received an update of their cumulative costs until that period. Subjects had to minimize cumulative travel cost over the duration of twenty-four periods. They were graded for a class assignment according to their performance; those with the lowest travel expenditures obtained the highest grades.[2] At the end of game they each filled out a questionnaire on cues and rules they had used in making their decisions.

We designed a 2 (quality information) $\times$ 2 (market share information) $\times$ 24 (time periods) experiment to run on the game. We manipulated the information levels between subjects, and the time factor within subjects. Quality information had two levels: subjects either had or did not have the airlines' historical on-time performance in each period. We did this to mimic typical segments that exist in the market based on information available to consumers. The market share information condition had two levels: subjects either were given or not given the market share of the airlines in the previous period. We did this to examine the effects of a dynamic market on market efficiency, in particular to test for the 'smart few' mechanism.

The range of prices of the airlines as well as their on-time performances was chosen to reflect actual prices and performances for airlines flying between LA and London (*circa* 1994 when the experiment was run). Local travel agents indicated that the price range was between $1000 and $1300 for a round-trip ticket. Newspaper reports revealed that the range of on-time arrival rates for airlines was 55–90%. We used these figures to determine the average prices and on-time performances of the ten airlines used in the experiment. Each airline had an average price and on-time performance within this range for the entire duration of the game. From period to period, the price and on-time performances of the airlines varied by $15 and 2%, respectively, around the mean. The price and on-time performance for each airline was chosen so as to achieve an overall price–performance correlation that fit the ambient price–quality correlation of airlines between LA and London.[3] The correlation was 0.44; specifically, on-time performance improved as price increased. However the correlation was not perfect; therefore, the best

[1] The airlines were given numbers as names to preclude unnecessary demand effects.
[2] Specifically, grades were assigned on a sliding scale going from 5 out of 10 for the highest expenditures to 10 out of 10 for the lowest.
[3] This was estimated from actual prices and on-time performances for real-world airlines using price data from travel agents and on-time rates published in newspapers.

Exhibit 3. Comparison of airlines on various attributes

Airline rank	1[a]	2	3	4	5	6	7	8	9	10
Expected value	1205	1206	1236	1273	1340	1377	1382	1413	1434	1451
Average price	1060	1090	1120	1040	1020	1290	1260	1180	1230	1160
Average quality (on-time performance)	80%[b]	83%	83%	67%	55%	88%	83%	67%	71%	59%
Comment	Best-value airline				Low-price airline	Premium airline				Worst brand

[a]Airlines are ranked in descending order of expected value.
[b]Represents the percentage of times the airline is on time.

value brand, defined as the one with the lowest expected value,[4] was neither the airline with the lowest price nor the one with best on-time performance. Exhibit 3 provides a comparison of the airlines chosen for the experiment on price-quality and value. For expository purposes the airlines are displayed in descending order of value.[5] Thus airline 1 and airline 10 were the best and worst brands respectively, defined in terms of expected value. Further, airline 6 had the best on-time performance as well as the highest price, while airline 5 had the worst on-time performance as well as the lowest price.

Subjects and procedure

Ninety seniors in an undergraduate program in business administration participated in the experiment. We ran each cell of the experiment as a different session in a large computer room. We randomly assigned subjects to each of these sessions, and within each session to a particular terminal. Further, we provided them with a cover story and detailed instructions, and allowed them two trial periods before the game began. We manipulated market-share information by giving feedback on other subjects' actions. In the condition in which this information was absent subjects made decisions independently of others' actions. In the condition in which this information was present, subjects made decisions with knowledge of other subjects' actions. Specifically in the latter condition, at the end of each period, all subjects in the market received information on the market share of each airline in that period. To do so, we tallied the number of consumers choosing each airline in the previous period and displayed this information on the board for all to see.

We encouraged subjects to deliberate as much as possible while making their decisions. After subjects entered their choice in a period, the computer program determined whether the flight was on time or late based on probabilities of delay. (These probabilities were the preset information on airline quality that may or may not have been provided to subjects depending on their condition.) Based on the cost of delay ($700), the computer calculated the total cost of choosing the airline for that period and displayed this information as feedback for the period.

At the end of 24 periods, the computer program provided subjects their final cumulative costs. The computer program also provided subjects with a questionnaire containing open- and closed-ended items designed to identify the cues and rules they used in making choices over time. Finally, we thanked and debriefed subjects.

[4] EV = ((Average Price) + ((1 − Average On-time Performance) × Cost of Delay)).
[5] This was not the order in which subjects were presented with the airlines.

RESULTS

We analyzed the purchase histories of individual subjects to (1) develop segments based on individual learning over the 24 periods, and (2) study the impact of learning at the individual and segment levels on market outcomes. In addition, we analyzed subjects' choices and their protocols to better understand the signals and rules they used in making choices over time. We especially focused on identifying subjects' whose choices remained suboptimal even with experience.

Developing segments

We designed an algorithm to segment subjects by their learning over time. We estimated learning of an individual by the trend of the individuals' cost of error in choices over time. Cost of error was the difference between the value of the chosen brand and the Best Brand in a period. A brand's value is the price of the brand plus the cost of delay. We used the following simple power function for this purpose:

$$Y_{ij} = a_0 X^{a1} \tag{1}$$

where Y_{ij} denotes the cost of error of subject i in period (1 to 20), X is one of the 20 periods of the experiment, and a_0 and a_1 are parameters to be estimated.

One advantage of this model is its flexibility; it takes on a variety of shapes depending on the value of a_1 (Lilien, Kotler and Murthy, 1992). If $a_1 = 0$, the function is a horizontal line, indicating no significant change in error costs with time. If $a_1 = 1$, there is a direct linear relationship: errors increase at a uniform rate, indicating people who grew steadily confused. If $a_1 < 0$, there is an inverse nonlinear relationship: errors decrease quickly and then level off, indicating early learning. Finally, if $a_1 > 1$, there is a direct nonlinear relationship: errors increase slowly at first and then rapidly, indicating people who grew confused later on. Another advantage of the model is that it may be estimated with the following linear regression model.

$$\log(ErrCost_j) = \log(a_0) + a_1 \log(period_j) + \varepsilon \tag{2}$$

We classified subjects into one of the four segments of Exhibit 1, by the following values of the parameters of the model:

(1) Nonlearners: a_1 not significantly different from zero and $a_0 > -4$,[6]
(2) Learners: a_1 significantly < 0
(3) Confused subjects: a_1 significantly > 1
(4) Early learners: a_1 not significantly different from zero and $a_0 < -4$.

The results showed that a large segment of subjects ($N = 55$ or 61.8%) did not learn, i.e. they made suboptimal choices throughout with no significant trend towards or away from optimality. The segment of early learners, i.e. those who made optimal choices throughout, was the second largest ($N = 20$ or 22.5%), followed by the segment of confused learners, i.e. those with a significant trend away from optimality ($N = 8$ or 9.0%). The segment of learners (i.e. those with a significant trend towards optimal decisions) was the smallest of the four segments ($N = 6$ or 6.7%). The histogram of values of a_1 for all subjects provides additional evidence in support of these results (see Exhibit 4).

[6] The intercept a_0 gives the initial error and is therefore used to distinguish between those who learned early ($a_0 < -4$) from those who did not learn ($a_0 > -4$). The value of −4 was determined empirically by comparing graphs with values of a_0 for individual cases.

Midpoint		Freq	Cum. Freq	Percent	Cum. Percent
-2	********	8	8	8.99	8.99
-1	*************************	25	33	28.09	37.08
0	***	45	78	50.56	87.64
1	*****	5	83	5.62	93.26
2	***	3	86	3.37	96.63
3	**	2	88	2.25	98.88
4	*	1	89	1.12	100.00

Exhibit 4. Histogram of parameter a_1 in equation (2)

The impact of individual learning on market learning

We traced the market shares of the best-value, the sixth-best and the worst-value airlines over the 24 periods in each segment as well as in the aggregate. In the segment of confused learners, the market share of the best-value airline was initially the highest, but in later periods it was the same as that of the other two airlines. In the segment of nonlearners, the market share of each of these three airlines remained about the same throughout (the market share of the best-value airline was, however, typically higher than that of the other two airlines). In the segment of early learners, the market share of the best-value airline increased gradually with time and was consistently higher than the share of the other two airlines. Finally, in the segment of learners, the market shares of the three airlines did not differ during the initial periods. In the later periods, however, the best-value airline emerged as clearly more popular than the other two.

Due to the preponderance of nonlearners, the market as a whole did not reflect any learning over time. For example, the relative market share of the three brands remained more or less constant over time. Further, the average error cost for the segment of nonlearners and the market over 20 periods was approximately flat (Exhibit 5).

We also tested our hybrid mechanism and Camerer's three mechanisms by which markets may be efficient when individual agents are not: the learning, cancellation, and smart-few mechanisms.

The learning hypothesis was not supported. Although a segment of learners was identified, the size of this segment was small relative to the market. There was little or no learning over time at the market level. Further, although subjects in one condition were given market-share information that signaled quality, this did not seem to have helped their learning.

The cancellation mechanism was not supported. The good decisions of learners were negated over time by the bad decisions of confused learners; as a result the cumulative cost of error of the entire market did not improve with time, but remained more or less constant. Further, the pattern of cost of error over time for the entire market was remarkably similar to that of the cost of error for the segment of non-learners. The net effect was that the market reflected the behavior of those who did not learn.

The smart few hypothesis was not supported. The 'learners' were too few compared to the 'nonlearners' and the 'confused learners' to make any impact on the market. This happened despite market-share feedback, a manipulation specially intended to let the decisions of learners influence the market.

Finally, the hybrid hypothesis was supported. Although a large number of subjects did not learn, very few made highly suboptimal decisions. For example, Exhibit 6 shows that the market share of the

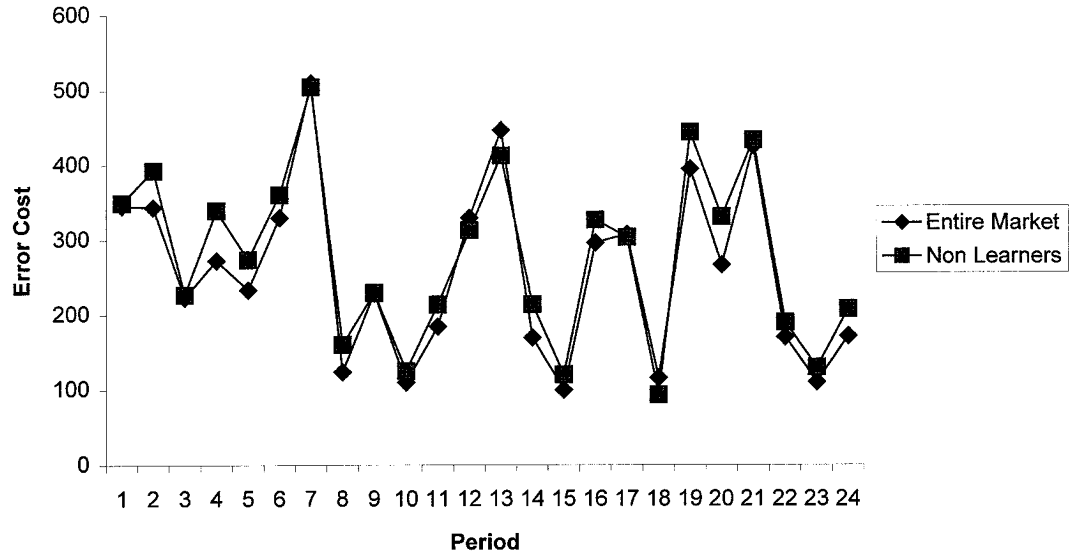

Exhibit 5. Error cost for segment of nonlearners versus the entire market over time

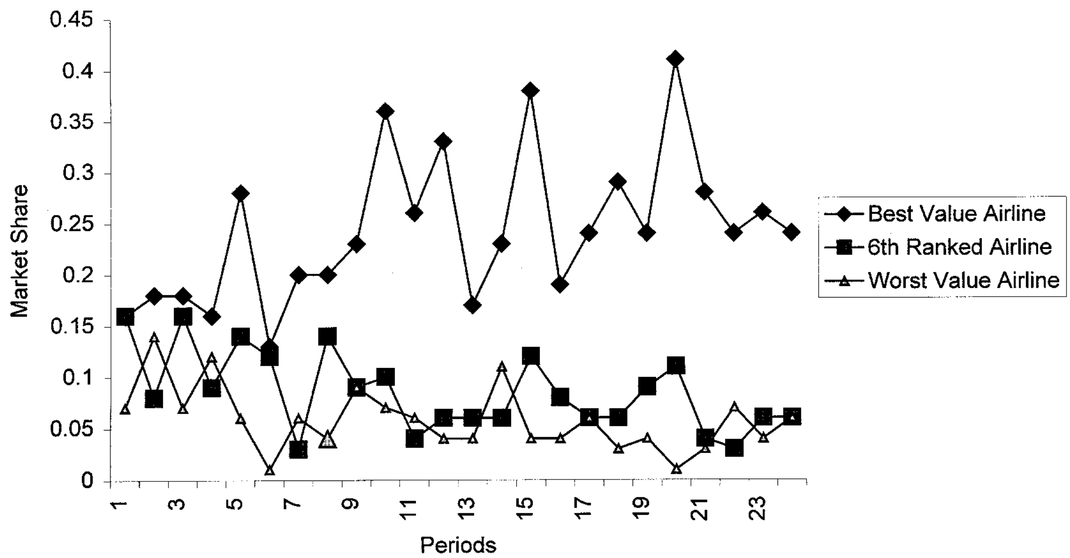

Exhibit 6. Market shares of best-value, sixth-ranked and worst-value airlines over time

sixth- and tenth-ranked airlines was consistently between 5% and 10%, settling down to around 5% by the last few periods. If the minimum market share necessary for an airline to survive is 5%, the tenth-ranked airline would be eliminated as early as the sixth period and market efficiency would automatically improve. However, if brands can survive with even a 1% market share, as often happens in real markets, all the airlines in the experimental market would survive and the market would not become efficient with time.

Exhibit 7. Airline choice percentages over 24 periods across conditions

Condition/airline value	Top-value airlines (airlines ranked 1, 2 and 3)*	Mid-range airlines (airlines ranked 4, 5, 6 and 7)	Worst value airlines (airlines ranked 8, 9 and 10)
No information	26.39	38.66	34.96
Market share only	56.12	18.24	25.66
Quality only	72.92	26.74	0.35
Quality and market share	80.21	18.89	0.89

*Airlines are ranked in descending order of expected value. Airline 1 and airline 10 are the best-value and worst-value airlines respectively.

Use of cues and heuristics in brand choice

To better understand subjects' learning processes, we carried out an analysis of subjects' brand choices over the duration of the game. The analysis revealed clear inefficiencies and irrationalities in the No Information and Market Share Only conditions. In particular, subjects in these two conditions chose the worst airlines almost as frequently as the best (see rows 1 and 2 of Exhibit 7). From subjects' responses to the questionnaire after the game, we found that these irrationalities arose because subjects used suboptimal rules in making choices. The two most frequently used rules in the No Information condition were 'choose neither a high-price nor a low-price airline' and 'if airline is on time re-choose, else switch' (used by 39% and 28% of the subjects respectively). These rules were also the most frequently used ones in the Market Share Only condition (32% and 22% respectively).

The use of the rule 'choose neither a high-price nor a low-price brand' resulted in the high choice frequency of the worst airlines (35% and 26% for the No Information and Market Share Only conditions respectively), as these brands all had average prices (see Exhibit 3). Furthermore, the protocols revealed that in the absence of information on quality, subjects viewed price as the most important signal of quality and felt that it was best to choose airlines with average prices (because high-priced airlines were too expensive and low-priced airlines might have poor quality). Subjects were not aware, however, that the use of this rule resulted in systematically suboptimal choices over time.

The analysis of brand choice reveals that subjects did use experience, but in a very short-sighted way. They used the rule 'if airline is on time re-choose, else switch'. This is equivalent to the well-known managerial maxim 'If it ain't broke, do not fix it'. Since being on time is a probabilistic event, even the best airlines can be late, and even the worst may be on time. So judgement based only on the most recent experience would not lead to enlightened learning from experience and the optimal choice.

The analysis of brand choice revealed fewer irrationalities and inefficiencies in the Quality Only and Quality and Market Share conditions. For instance, the choice of worst-ranked airlines dropped to less than 1%. However, the analysis revealed one outstanding example of irrationality. Specifically airline 6, despite being dominated by five other airlines, was chosen with high frequency in both conditions (18% and 8% respectively). Once again this was due to subjects' use of heuristics. In both conditions the most frequently used heuristic was 'choose a high-quality airline' (33% and 57% respectively). As airline 6 was the highest-quality airline, it was chosen frequently, resulting in inefficiencies over time. The use of this rule also resulted in the lowest quality airline, airline 5, being chosen less frequently than airline 6 in the Quality Only (2% versus 18%) and Quality and Market Share conditions (5% versus 8%), even though airline 5 had a better expected value than airline 6. The protocols revealed that in these two conditions, because subjects were provided information on quality, they relied on it heavily in making choices over time. This in turn helped them avoid the worst-value airlines, but led them into other systematic irrationalities.

To formally test subjects' use of cues and rules, we performed a logit analysis of subjects' brand choices on product attributes and experimental market conditions. Specifically, our model was of the form:

$$Y = \text{logit}(p) = \log(p/(1-p)) = \alpha + \beta' X \quad (3)$$

where X is the matrix of explanatory variables, Y is the vector of transformed choices, α is the intercept parameter, and β is the vector of slope parameters (Train, 1990). The dependent variable, choice of airline, was 1 if the airline was chosen in a particular period, 0 if it wasn't. The main independent variables were:

Price:	the price of the airline
Quality × QInformation:	on-time performance × whether that information was provided
Market share information:	last period's brand share × whether that information was provided
Last Success:	whether the airline chosen in the last period was on time or delayed.

Because subjects often used high or low quality and high or low price as signals of good value, we also included dummy variables indicating whether airlines were the highest or lowest price or quality in each choice situation.

The results provided several insights about consumer choice behavior (see Exhibit 8). First, the results show that brands with lower price had a higher probability of choice. This is the only response that is normatively reasonable. All the other significant effects indicate cue-based choice. Second, quality × Q information is not significant, indicating that consumers who had quality information did not assess it proportionately. Rather, highest quality and lowest quality are significant, indicating that subjects favored the brand with the highest quality and avoided the one with the lowest. Normatively optimal behavior would have been for subjects to integrate quality proportionately with the probability and cost of being late. Third, the single most important variable in the model is last successful

Exhibit 8. Results of logistic regression for airline choice

Independent variable	*B* coefficient	Chi-square
Price	−0.37[a]	126.05
Quality × Qinformation*	−0.003	0.05
Highest price**	0.04[e]	2.05
Lowest price**	−0.12[a]	28.53
High quality**	0.09[c]	7.52
Low quality**	−0.16[a]	69.68
Market-share information*	−0.03[d]	3.94
Last success**	0.52[a]	3283.55
Model AIC (Chi-Sq)	9826.62	4240.40

*Quality is on-time performance of airline; QInformation is information on quality present.
**Highest price is 1 for highest-priced airline, 0 otherwise;
Lowest price is 1 for lowest-priced airline, 0 otherwise;
Highest quality is 1 for highest-quality airline, 0 otherwise;
Lowest quality is 1 for lowest-quality airline, 0 otherwise;
Last success is 1 if airline was chosen *and* was on time, 0 otherwise;
*Market-share information is last period market share × market share provided to subjects.
[a] $p < 0.0001$.
[b] $p < 0.001$.
[c] $p < 0.01$.
[d] $p < 0.05$.
[e] $p < 0.10$.

experience. (This can be seen both by the size of the coefficient and the Chi-square statistic.) The implication is that consumers were primarily motivated by their last experience. That is, if their experience was successful they repurchased, otherwise they switched.

In sum, these results provide evidence that subjects' choices were driven by rules which especially used extreme quality levels and immediate past experience as signals of the value of the airlines.

DISCUSSION

This section first discusses the main conclusions, then points out some limitations of this study and closes with directions for future research.

Main conclusions

Several important findings emerge from the study. First, it confirms the general finding that consumers make suboptimal decisions when product quality is uncertain, resulting in inefficient markets. The main reason is that consumers do not understand uncertain quality and misuse information dealing with it, even over time. Second, the study shows that consumers do not all behave in the same way. Clearly, groups of consumers do learn from experience while others do process uncertainty appropriately. However, the market as a whole remains inefficient because of the preponderance of nonlearners. Furthermore, segments of consumers who learn wrong rules may be large enough to negate the benefits of the segment of informed decision makers, the 'smart few'. Thus Camerer's cancellation hypothesis is not supported.

Third, consumers make suboptimal choices because they resort to simple decision rules, rather than the optimal integration of price and quality information. For example, in the absence of quality information consumers use a simple rule such as 'choose only average priced brands', even though these brands may be dominated by cheaper or more expensive brands. Or, when given quality information, they might use a simple rule 'choose only high-quality brands' even though high-quality brands may not provide the best value. Most importantly, consumers overweigh recent experience, a practice which is captured in their use of the rule 'if airline is on time re-choose, else switch'. This finding is consistent with findings that outcome-feedback can be deleterious to learning (Remus, O'Connor and Griggs, 1996; Siegel-Jacobs and Yates, 1996; Tellis and Gaeth, 1990; Thompson and DeHarpport, 1994). Because feedback draws attention to the most recent observation, consumers are likely to overweigh recent experience in making subsequent decisions. As a result, consumers switch from a good-value airline with only one bad experience, or stick to a high-quality airline which is greatly overpriced, both of which lead to suboptimal decisions even with experience. This finding is the key explanation for our disconfirmation of Camerer's learning hypothesis. This finding is quite relevant because outcome-feedback is a feature of real world consumption of both search and experience goods.

Fourth, the study provides only limited support for the smart few mechanism. Specifically, adding market-share information (thus making the market responsive to feedback) improved the quality of choices a little, but irrationalities continued to exist. Presumably, therefore, the experience of a few smart consumers does influence the choices of the rest, but not to the extent that the market becomes completely efficient.

Finally, the study supports the hybrid hypothesis. Inefficient brands tend to have lower market share than efficient brands. Thus if the break-even market share is relatively low, inefficient brands may still survive due to the persistent suboptimality of a large segment of consumers and the simplistic decision rules they use.

Limitations

A possible weakness of all experiments is their external validity. It is possible that consumers in real markets, with real experience and motivation, are more likely to learn than under experimental conditions. However, three features of real markets — greater complexity, less motivation, and less timely feedback — could make learning in real markets more difficult and unlikely than in experiments. First, in real markets, consumers choose simultaneously from different categories; in this study choice was restricted to a single category. This makes the learning task in real markets considerably more complex than the task in this experiment. Second, in some markets (as in low-involvement product categories such as paper towels) consumers may not be as motivated to exercise rationality with the same stringency as in this experiment. In fact, conversations with the subjects indicated that they were adequately motivated but in many cases could not identify the best brand. Thus, satisficing may result in greater tolerance of suboptimality than decision theory allows. Finally, feedback in real markets is likely to be less timely than in this experiment. Without feedback there can be no learning. All these three factors suggest that learning is likely to be less widespread in real markets than in experiments. This prognosis, however, is not altogether bleak. Although irrationality exists, it is not so egregious that highly inferior brands survive. Indeed the hypothesis that highly inferior brands will not receive the market share necessary to stay viable indicates that even with individual irrationality markets may stay inefficient but not grossly so.

This conclusion is supported with evidence from field data on two scores. First, there is evidence from research on scanner panel data that the strongest predictor of choice in a product category is past choice as indicated by the parameter of a loyalty variable (Guadagni and Little, 1983). This suggests that even when consumers are presented with information on price and quality that helps them to make rational choices, they stick with brands that they are familiar with even if these brands are not the best choice. Second, there is evidence that the correlation between price and quality in over 1200 durable and non durable consumer goods categories is far from perfect and sometimes even negative (Tellis and Wernerfelt, 1987). These authors suggest that imperfect information about quality is a major cause of this situation. Thus real markets stay relatively inefficient even in equilibrium.

Future research

Several limitations of this paper suggest areas for further research on consumer learning. First, the experiment was designed to test for learning in the case when supply is static and demand is evolving. In the real world, however, consumers are further challenged by rapidly evolving product categories with new brands and innovations continually entering the marketplace. How do consumers learn when supply too is evolving? What cues do they now use in making choices, and how do the use of these cues and rules evolve with experience?

Second, the experiment was designed such that consumers all entered the market at the same time, i.e. at period 1 all subjects were equally unfamiliar with the product category. In the real world, however, consumers enter markets at different times; in addition to the ability or motivation to learn, consumers will differ at a given time on the basis of their experience with the category. The effect of disaggregate behavior on market efficiency will be further complicated by the heterogeneity in consumer experience. This offers a challenging task to those interested in dynamically modeling markets from the bottom up.

Finally, while the market employed in the study provided some conditions necessary for testing feedback-based learning (such as providing period to period market share information), it did not allow for market-selection mechanisms which might otherwise ensure that real markets are efficient even if individuals are not. Thus, future research may extend the results of this study by designing experiments in which only rational consumers (and, by extension, efficient brands) are allowed to survive and all others are forced out of the market over time. Such experiments may shed still more light on the important and interesting issue of individual rationality over time and its relationship to market efficiency.

ACKNOWLEDGEMENTS

The authors would like to thank Gary Gaeth, Avu Sankaralingam, David Stewart, three anonymous reviewers and the editor for their useful comments.

REFERENCES

Alchian, A. 'Uncertainty, evolution, and economic theory', *Journal of Political Economy*, June (1950), 211–221.
Bettman, J. R., Johnson, E. J. and Payne, J. W. 'Consumer decision making', in Robertson, T. S. and Kassarjian, H. H. (eds), *Handbook of Consumer Behavior*, Englewood Cliffs, NJ: Prentice Hall, 1991.
Camerer, C. F. 'The rationality and prices and volume in experimental markets', *Organizational Behavior and Human Decision Processes*, **51** (1992), 237–272.
Einhorn, H. J. and Hogarth, R. M. 'Confidence in judgment: persistence in the illusion of validity', *Psychological Review*, **85** (1978), 395–416.
Guadagni, P. and Little, J. D. C. 'A logit model of brand choice calibrated on scanner data', *Marketing Science*, **2** (1983), 203–238.
Haines, G. *Consumer Behavior: Learning Models of Purchasing*, New York: The Free Press, 1969.
Hoch, S. J. and Ha, Y. 'Consumer learning: advertising and the ambiguity of product experience', *Journal of Consumer Research*, **13** (1988), 221–233.
Kahneman, D. and Tversky, A. 'Prospect theory: an analysis of decision under risk', *Econometrica*, **47** (1979), 263–91.
Kahneman, D., Knetsch, J. L. and Thaler, R. H. 'Fairness and the assumptions of economics', in Hogarth, R. M. and Reder, M. W. (eds), *Rational Choice: The Contrast Between Economics and Psychology*, Chicago and London: University of Chicago Press, 1987.
Kennedy, P. W. 'Word-of-mouth communication and price as a signal of quality', *Economic-Record*, **70** (1994), 373–80.
Lilien, G., Kotler, P. and Murthy, K. S. *Marketing Models*, Englewood Cliffs, NJ: Prentice Hall, 1992.
Lucas, R. E. Jr. 'Adaptive behavior and economic theory', in Hogarth, R. M. and Reder, M. W. (eds), *Rational Choice: The Contrast Between Economics and Psychology*, Chicago and London: University of Chicago Press, 1987.
Milgrom, P. and Roberts, J. 'Price and advertising signals of product quality', *Journal of Political Economy*, **94** (1986), 796–821.
Miller, R. M. and Plott, C. R. 'Product quality signaling in experimental markets', *Econometrica*, **53** (1985), 837–872.
Monroe, K. B. and Petroshius, S. M. 'Buyers' perceptions of price: an update of the evidence', in Kassarjian, H. H. and Robertson, T. S. (eds), *Perspective in Consumer Behavior*, Glenview, ILL: Scott, Foresman and Company, 1981.
Remus, W., O'Connor, M. and Griggs, K. 'Does feedback improve the accuracy of recurrent judgmental forecasts?', *Organisational Behaviour and Human Decision Processes*, **66**(1) (1996), 22–30.
Salop, S. and Stiglitz, J. 'Bargains and rip-offs: a model of monopolistically competitive price dispersion', *Review of Economic Studies*, **44** (1977), 493–510.
Siegel-Jacobs, K. and Yates, J. F. 'Effects of procedural and outcome accountability on judgement quality', *Organisational Behaviour and Human Decision Processes*, **65**(1) (1996), 1–17.
Smallwood, D. E. and Conlisk, J. 'Product quality in markets where consumers are imperfectly informed', *Quarterly Journal of Economics*, **43** (1979), 1–23.
Spence, A. M. *Market Signaling: Informational Transfer in Hiring and Related Screening Processes*, Cambridge, MA: Harvard University Press, 1974.
Stiglitz, J. E. 'The theory of screening, education and the distribution of income', *American Economic Review*, **65** (1975), 283–300.

Tellis, G. J. and Fornell, C. 'Advertising and quality over the product life cycle: a contingency theory', *Journal of Marketing Research*, **15**(1) February (1988), 64–71.
Tellis, G. J. and Gaeth, G. J. 'Best value, price-seeking, and price aversion: the impact of information and learning on consumer choices', *Journal of Marketing*, **54** (1990), 34–45.
Tellis, G. J. and Wernerfelt, B. 'Competitive price and quality under asymmetric information', *Marketing Science*, **6**(3) (1987), 240–253.
Thompson, L. and DeHarpport, T. 'Social judgement, feedback, and interpersonal learning in negotiation', *Organisational Behavior and Human Decision Processes*, **58**(3) (1994), 327–345.
Train, K. *Qualitative Choice Analysis: Theory, Econometrics and an Application to Automobile Demand*, Cambridge, MA: The MIT Press, 1990.
Wolinsky, A. 'Prices as signals of product quality', *The Review of Economic Studies*, **50** (1983), 647–58.
Tversky and Kahneman (1986).

Authors' biographies:
Jaideep Prabhu is Assistant Professor of Marketing, Tilburg University, The Netherlands.

Gerard J. Tellis is the Neely Professor of Marketing at the Marshall School of Business, University of Southern California, USA.

