PETER S. FADER and LEIGH McALISTER*

Because consumers are believed to use simplifying heuristics to screen alternatives before choosing one to purchase, the authors formulate and estimate an elimination by aspects model at the individual household level, using UPC scanner data collected from purchases of regular ground coffee. The proposed model fits and predicts the data as well as its compensatory approximation. The parameters of the proposed model are reasonably stable, unbiased, and different in interesting ways from their compensatory counterparts.

An Elimination by Aspects Model of Consumer Response to Promotion Calibrated on UPC Scanner Data

A busy consumer is pushing her shopping cart down one of the aisles of her favorite supermarket. As she nears the coffee area she scans the brands, looking for special offers. She notices an acceptable brand on promotion, puts a can of that promoted coffee in her cart, and moves down the aisle.

In today's supermarkets, where different promotional devices assail consumers from all directions, this type of behavior is quite plausible. Consumer promotions, often offering substantial price cuts, enable shoppers to make quick, fairly good decisions without processing all available information. Brand choice models rarely allow promotions to play this type of role. The model proposed here is based on the notion of a phased decision strategy. We hypothesize that, under certain conditions, some individuals screen alternatives for promotional offers before making a selection.

In recent years, the notion of a phased, or noncompensatory, decision strategy has gained wide acceptance. Under a variety of conditions, consumers are believed to use simplifying heuristics to restrict their choice set to a limited number of alternatives before choosing one to purchase. As Johnson, Meyer, and Ghose (1989, p. 255) have stated,

> Even the most casual examination of consumers' verbal reports indicates that decisions are based, at least in part, on *noncompensatory* strategies that do not involve tradeoffs. Supporting evidence for the use of such heuristics pervades the literature, including studies using information display boards (e.g., Bettman and Jacoby 1976, Lussier and Olshavsky 1979, Payne 1976), eye movements (Russo and Dosher 1983), and even in-store verbal protocols (Payne and Ragsdale 1978).

Johnson, Meyer, and Ghose call for models that are "better representations of decision processes" (p. 268). We respond to that call by offering an elimination by aspects (EBA) model of consumer choice that can be estimated with UPC scanner data. We find, as would be predicted (Dawes and Corrigan 1974; Johnson and Meyer 1984), that these individual-level EBA models perform no better than their compensatory counterparts in estimation and prediction. However, these EBA models reflect a widely accepted noncompensatory functional form, yielding parameters that are unbiased and reasonably stable. Further, those parameters provide insights that would be masked by a compensatory modeling scheme.

We first present some background about EBA and related models and then describe the model used in our application. In the Estimation section we describe the

*Peter S. Fader is Assistant Professor, Department of Marketing, University of Pennsylvania. Leigh McAlister is Associate Professor, College of Business Administration, University of Texas, Austin.

The authors thank Carnation, Procter & Gamble, and the Marketing Science Institute for providing financial support for the research; Selling Areas Marketing, Inc. (SAMI), Peter Guadagni, and John D. C. Little for providing data; and many colleagues in academia and industry for useful comments.

Journal of Marketing Research
Vol. XXVII (August 1990), 322–32

data, estimation technique, and reference models and report results of statistical tests of parameter stability. We then consider the differences between noncompensatory and compensatory models' parameters. We close with the argument that this model complements current models and estimation techniques and that it provides insights that are otherwise masked.

THE ELIMINATION BY ASPECTS MODEL

The EBA model, first described by Tversky (1972), is perhaps the best-known model of sequential, noncompensatory choice behavior. The basic premise behind EBA is that the decision maker chooses an attribute, screens out unacceptable alternatives on the basis of that attribute, and continues until the remaining alternatives do not share any common attributes. At that point the Luce choice axiom (Luce 1959) is used to choose a single alternative from those that remain.

Despite the intuitive appeal of EBA, applications of the model have been rare. Batsell and Polking (1985) cite two primary reasons: (1) the full EBA model requires a large number of parameters to be estimated and (2) software designed specifically for EBA estimation is lacking.

Several researchers have proposed extensions or variations of EBA (see, e.g., Batsell and Polking 1985; Manrai and Sinha 1989), whereas others have proposed alternative noncompensatory models with no direct links to EBA (e.g., Currim, Meyer, and Le 1988; Gensch 1987). Most of their studies have successfully uncovered the frequent use of phased heuristics, demonstrating that the explicit estimation of such decision strategies can significantly improve a model's fit and interpretability. Here, however, we examine the consequences of using a "pure" EBA model with no further analytical adornments. We believe that an elementary EBA model, featuring a simple screening rule, can contribute to our understanding of individuals' choice behavior.

MODEL DESCRIPTION

The supermarket is a fertile environment for the use of simplifying heuristics, because consumers make hundreds of low involvement decisions on a regular basis. It seems reasonable that consumers might often use a limited set of attributes to screen out "unacceptable" alternatives, thereby reducing the complexity of the decisions they must make. A prominent attribute in this situation is promotion. For example, the CLS estimation procedure used by Currim, Meyer, and Le (1988) suggests that promotions such as in-store displays and newspaper features play a prominent role in the phased decision strategies apparently used by numerous consumers in purchasing regular ground coffee. We highlight and further explore the role of promotions as simple screening devices in low involvement shopping situations.

We begin with an assumption about consumer preferences.

A_1: Each consumer has adopted, through past behavior, a set of brands that he or she will consider in purchasing from a product category. We call that collection of brands the "acceptable set."

We next make a series of assumptions about the consumer's choice rule and the role of promotional offers in that rule.

A_2: The consumer chooses among brands considered for choice with probabilities proportional to his or her preferences for those brands.

A_3: The consumer's promotion sensitivity is defined as the probability that he or she will screen alternatives on the basis of promotional offers that he or she encounters.

A_4: When screening on promotional offers, the consumer restricts consideration to those brands in his or her acceptable set that are on promotion.

A_2 is consistent with Luce's (1959) formulation and has been shown to have strong predictive validity across a wide variety of choice situations. The formulation is less successful when some alternatives in the choice set are more similar to one another than they are to other alternatives (e.g., two caffeinated coffees are more similar to one another than is a caffeinated coffee to a decaffeinated coffee.) We consider only choice from product categories in which such unequal similarity is not an important characteristic of choice alternatives.[1]

Just as A_2 reflects stochasticity in the brand choice, A_3 reflects stochasticity in promotion response. Here we assume that a consumer occasionally will miss or intentionally ignore promotional offers. Combining A_3 with A_4 suggests that the consumer sometimes will react to promotional offers by screening out unpromoted brands.

Definition of Terms

As detailed in A_1, we assume that each consumer has a set of acceptable brands, denoted a_i. Each consumer also has a promotion sensitivity parameter, γ_i, which represents the probability that consumer i will restrict his or her choice to promoted brands, conditional on at least one of the acceptable brands being on promotion. Finally, each consumer has a set of preferences for the brands in a_i. Consumer i's preference for brand k, π_{ik}, is a number between 0 and 1 that reflects his or her feelings toward all features of brand k other than its promotional status. Let

a_i = set of acceptable brands for consumer i (those brands that consumer i ever chooses),

π_{ik} = consumer i's preference for brand k, independent of brand k's promotional status, $0 \leq \pi_{ik}$, $\Sigma\pi_{ik} = 1$,

[1]This restriction of equal similarity across choice alternatives is not necessary in a completely general EBA model. We make the restriction here to minimize the number of parameters that must be estimated for each individual-level EBA model.

γ_i = consumer i's promotion sensitivity (the probability that consumer i will restrict his or her choice to promoted acceptable brands given that some acceptable brand is on promotion),

$$P_{kt} = \begin{cases} 1 \text{ if brand } k \text{ is on promotion at time } t \\ 0 \text{ if brand } k \text{ is not on promotion at time } t, \text{ and} \end{cases}$$

E_{ikt} = probability, assigned by the EBA model, that consumer i will choose brand k at time t.

The EBA Model

The analytical EBA model focuses on the role of promotion in a phased decision strategy. Given that some acceptable brand is on promotion, consumer i decides first whether to screen for promotion and then makes a brand selection from the resulting relevant set. This process is denoted as in the following equations.

If some acceptable brand is on promotion:

$$\text{(1a)} \qquad E_{ikt} = \underbrace{(\gamma_i)}_{\text{I}} \underbrace{\frac{\pi_{ik} P_{kt}}{\sum_{m \epsilon a_i} \pi_{im} P_{mt}}}_{\text{II}} + \underbrace{(1 - \gamma_i)}_{\text{III}} \underbrace{\frac{\pi_{ik}}{\sum_{m \epsilon a_i} \pi_{im}}}_{\text{IV}} .$$

If no acceptable brand is on promotion:

$$\text{(1b)} \qquad E_{ikt} = \frac{\pi_{ik}}{\sum_{m \epsilon a_i} \pi_{im}}$$

If some acceptable brand is promoted at time t, consumer i will screen on promotion with probability γ_i and will choose from the entire acceptable set with probability $1 - \gamma_i$. If screening on promotion, consumer i will choose brand k with the probability defined in term II of equation 1a. If not screening on promotion, consumer i will choose brand k with the probability defined in term IV. Terms II and IV imply that consumer i will choose among relevant brands with probabilities proportional to preferences.

If no acceptable brand is on promotion at time t (i.e., $P_{mt} = 0$ for each brand m in the consumer's acceptable set), the promotion screening process does not apply; therefore, the EBA choice probability expression, E_{ikt}, collapses to equation 1b.

For the proposed EBA model, the only shared attribute is promotion; the consumer can either use it (with probability γ_i) to screen out unpromoted brands or not use it. In either case, the Luce axiom is used to choose among the remaining alternatives. Given the assumptions of our EBA model (i.e., $\gamma_i \epsilon$ [0,1], $\Sigma\pi_{ik} = 1$), our EBA choice probability expression (1a and 1b) can be shown to be consistent with Tversky's (1972) more general model of phased decision making.

Though the structure and assumptions of our EBA model are likely to be a fair basis for many low involvement brand choice decisions, they also may be somewhat restrictive. In particular, one might want to include other predictor variables. We limit this model to brand preferences and promotion sensitivities so that estimation can be done at the individual level. A more complex model incorporating regular price, depth of promotional price cut, and other marketing variables would jeopardize the quality of parameter estimates given the shortage of data at the individual level. We justify our choice of variables to include by noting that in Guadagni and Little's (1983) study of purchase behavior, the most important predictors were brand "loyalties" (an analogue to brand preferences) and a dummy variable registering the presence or absence of promotion. (Currim, Meyer, and Le 1988 offer similar findings.) Furthermore, there is reason to doubt the criticality of incorporating regular price and depth of promotional price cut. Scanner data show there is very little fluctuation in the regular price of most brands. In addition, there is very little fluctuation in the depth of promotional price cuts for brands in a given product category. We offer these facts as justification of our model specification. The ability of the individual-level EBA models to fit and predict choices better than an aggregate model incorporating those additional predictors adds empirical support for our formulation. In future research, many of the assumptions and omitted variables can be considered as elaborations of the basic model.

ESTIMATION

Investigations into phased decision making often have been based on verbal protocol data. For such analyses, consumers are asked to "think aloud" while making a decision. Recordings of the protocols are later transcribed and analyzed. These studies have provided many insights unattainable with techniques such as conjoint analysis and regression analysis as they traditionally have been applied. However, protocol analysis has been criticized on at least two dimensions (Bettman 1979). First, protocol analysis requires a great amount of data that is time consuming to collect. This data constraint has driven researchers to very small sample sizes (e.g., Bettman 1970 investigated two subjects). Second, the quality of protocol data has been questioned. Consumers may censor their reports, they may be unable to verbalize what they actually are doing or, when describing a low involvement decision that they have made many times in the past, they may be producing a retrospective rationale.

Some of these drawbacks can be overcome by using other information acquisition methods. Popular alternatives include information display boards (or software), eye movement analysis, and chronometric analysis (Bettman, Johnson, and Payne 1989). These techniques can provide more data than protocol analysis, but each has idiosyncratic drawbacks that might limit its generalizability to actual choice situations.

We propose a complementary technique for investigating phased decision-making behavior that exploits the rich and accurate data now available from universal product code (UPC) scanner panels. Though this technique does not provide the richness of insights possible with consumer self-reports, it does allow the investigation of a fairly simple, phased decision process that might be expressed in repetitive, low involvement choice. It also allows investigation of many different consumers' behavior over a long time span.

We estimate the EBA models by using a general maximum likelihood routine and UPC scanner panel data for 200 consumers prepared by Guadagni and Little (1983). These data, collected at supermarket checkout counters, create a detailed record of each purchase occasion.[2] Scanner data also contain comprehensive information about price, availability, and promotional activity for each brand for every week.

A likelihood function is obtained for consumer i by using equation 1 for the probability of consumer i selecting each actually chosen brand and then multiplying these probabilities together. We used a FORTRAN program with nonlinear optimization subroutine to choose optimal values for γ_i and the set of π_{ik}.[3]

We next describe the reference models. The first one is an individual-level model and is based on a compensatory approximation of EBA. The second reference model is a widely cited aggregate compensatory approximation of EBA. We offer this second comparison to show that predictive power need not be lost by turning to individual-level models.

Reference Models

The individual-level compensatory approximation to EBA that we use is a disaggregate multinomial logit model. The aggregate compensatory approximation is Guadagni and Little's (1983) logit model. The models are fit by using the Guadagni and Little calibration sample panelists with the first 45 weeks of data. The last 20 weeks of data are held out for predictive testing.[4]

Disaggregate multinomial logit (DMNL). Like EBA, disaggregate multinomial logit reflects idiosyncratic preferences and idiosyncratic promotion sensitivities. DMNL can be thought of as a compensatory approximation of EBA.[5] The most common form of DMNL uses the following type of function form:

$$D_{ikt} = \frac{e^{\text{consumer } i\text{'s utility for brand } k \text{ at time } t}}{\sum_{m \in a_i} e^{\text{consumer } i\text{'s utility for brand } m \text{ at time } t}}$$

where "consumer i's utility for brand k at time t" is expressed as a linear function of brand k's attributes at time t. In the simple model considered here, "consumer i's utility for brand k at time t" is equal to

$$b_{i0} P_{kt} + \sum_{m \in a_i} b_{im} X_m$$

where P_{kt} indicates whether brand k is on promotion at time t and the X_m's are dummy variables indicating the brand under consideration (e.g., when estimating utility for brand k, $X_k = 1$ and all other $X_m = 0$). Hence, "i's utility for brand k at time t" $= b_{i0} P_{kt} + b_{ik}$ and

$$D_{ikt} = \frac{e^{b_{i0}P_{kt}+b_{ik}}}{\sum_{m \in a_i} e^{b_{i0}P_{mt}+b_{im}}}.$$

We restate DMNL to facilitate its comparison with EBA. In particular, we let

$e^{b_{i0}} = \Theta_i =$ consumer i's promotion sensitivity under DMNL

and

$e^{b_{ik}} = \Pi_{ik} =$ consumer i's preference for brand k under DMNL

(scaled so that $0 \le \Pi_{ik} \le 1$ and Π's sum to 1).

$$D_{ikt} = \frac{\Theta_i^{P_{kt}} \Pi_{ik}}{\sum_{m \in a_i} \Theta_i^{P_{kt}} \Pi_{ik}} \tag{2}$$

This restatement highlights the different roles promotion plays in the two models and facilitates the comparison of estimates of DMNL's preference parameters (Π_{ik}) with estimates of EBA's preference parameters (π_{ik}).

An example helps illustrate the structural difference between the two models. Suppose a consumer has two relevant brands ($a =$ [brand 1, brand 2]) with preferences π_1 and π_2, respectively, for EBA and Π_1 and Π_2, respectively, for DMNL. We drop the subscript i for expositional purposes.

If at time t only brand 2 is on promotion, then

$$D_{1t} = \frac{\Pi_1}{\Pi_1 + \Theta\Pi_2} \tag{3}$$

[2]To this point we have assumed implicitly that each decision maker in our database is an individual. We might be more correct to use the term "household," because each panelist identification number refers to a specific household rather than an individual. Researchers have emphasized and tried to account for differences between individual and household brand choice decisions (Davis 1976; Kahn, Morrison, and Wright 1986). These individual/household distinctions may have some influence on the estimated parameters, but should not affect the general theory to a large extent.

[3]Estimated parameter values were reasonably insensitive to starting points and to the version of the optimization routine employed.

[4]Each panelist's acceptable set is defined as the set of brands chosen at least once by the panelist over the 45 weeks of calibration data. Any brands first purchased during the 20 weeks of holdout data lead to choice probabilities of zero for EBA and DMNL.

[5]Under certain restrictive conditions, DMNL is a special case of EBA—for example, if we further restricted the proposed modeling context to consider only brand preferences.

and

$$D_{2t} = \frac{\Theta \Pi_2}{\Pi_1 + \Theta \Pi_2}. \tag{4}$$

For the EBA model we have

$$E_{1t} = (1 - \gamma) \frac{\pi_1}{\pi_1 + \pi_2} \tag{5}$$

and

$$E_{2t} = (\gamma) \frac{\pi_2}{\pi_2} + (1 - \gamma) \frac{\pi_2}{\pi_1 + \pi_2}. \tag{6}$$

Comparing the DMNL and EBA model structures, we see that they imply very different consumer choice strategies. In the DMNL model, utility for promoted brands is boosted by Θ and the consumer chooses from acceptable brands with probabilities proportional to the altered utilities. In the EBA model, promotion is used as a signal for screening alternatives. The two models differ in how often a consumer will respond to promotions. The EBA model suggests that $1 - \gamma$ of the time the consumer will deviate from his or her promotion screening strategy. The DMNL model allows for no deviations from its prescribed decision strategy; it assumes that utility is always boosted by the presence of promotions.

Guadagni and Little (GL). Guadagni and Little's (1983) aggregate logit model captures individual differences in preferences through their loyalty variables. It goes beyond those loyalty variables to reflect (among other things) promotion sensitivity. However, promotion sensitivity in their model is estimated at the market level. We refer the reader to their article for an exact specification of their model, but note here that their model has been shown to fit these data well and to have very good predictive ability.

Comparison Statistics

Table 1 reports statistics comparing the three models across the estimation and forecast periods. Estimation period statistics include $\overline{PC}$, the average probability assigned to the brands actually chosen; RMSE, root mean squared error; LL, the log of the likelihood function; and $\bar{\rho}^2$ (Horowitz 1983), an estimate of goodness of fit that adjusts LL for degrees of freedom.[6] Smaller values of RMSE indicate a better fit whereas larger values of all other statistics indicate a better fit.

Estimation period values of $\overline{PC}$, RMSE, and LL indicate that EBA and DMNL are indistinguishable on the basis of a simple t-test ($p > .10$) and that both fit the data better than the Guadagni and Little (GL) aggregate model ($p < .01$). However, EBA and DMNL estimate

Table 1
MODEL COMPARISON STATISTICS

	Estimation period (first 45 weeks)				Forecast period (last 20 weeks)		
Model	$\overline{PC}$	*RMSE*	*LL*	$\bar{\rho}^2$	$\overline{PC}$	*RMSE*	*% 1st choice predictions*
EBA	.670	.445	−595	.60	.545	.587	43.6
DMNL	.670	.447	−602	.59	.548	.588	43.4
GL	.531	.588	−977	.48	.494	.628	42.3

many more parameters than GL. The $\bar{\rho}^2$ statistic adjusts for degrees of freedom and still shows an advantage for EBA and DMNL.

Forecast period statistics include $\overline{PC}$, RMSE, and the percentage of correct first-choice predictions. The likelihood-based statistics (LL and $\bar{\rho}^2$) are omitted because they are undefined for consumers who chose some brand in the forecast period that they had not chosen in the estimation period.[7] The differences between the aggregate reference model (GL) and the two individual-level models (EBA and DMNL) are significant at the .01 level. The difference between EBA and DMNL is not significant at the .10 level.

We now focus on the two prediction schemes (EBA and DMNL) that best fit and predict the data. Because of the small sizes of the samples used to estimate these individual-level models (average number of purchases per consumer in the estimation period = 15), we cannot rely on standard asymptotic test statistics to evaluate their parameters. Therefore, we next describe and report the results of a bootstrap procedure used to examine the stability and unbiasedness of the individual-level models' parameters. We then characterize differences in parameters estimated by EBA and DMNL. The upshot of those differences is explored in the Implications section.

Quality of Parameters

A typical consumer in this database has three brands in the acceptable set and makes roughly 15 purchases during the estimation period. The individual-level models (EBA and DMNL) estimate four parameters for such a consumer (one promotion sensitivity parameter and three brand preference parameters). Given the small number of observations per parameter estimate, we seek some assurance that the estimated values of these parameters are not overly affected by the length of the purchase histories and the variability of the specific promotional environments faced by each consumer.

---

[6] $\bar{\rho}^2 = 1 - ([LL - (k/2)]/LL_0)$ where k is the number of estimated parameters, LL is the log likelihood of the estimated model, and LL_0 is the log likelihood of a reference model. In this case, we use aggregate market shares as the reference model.

[7] Recall that consumer i's acceptable set, a_i, is defined as any brand chosen by consumer i during the estimation period. If, in the forecast period, consumer i chooses some brand that is not in a_i, any of the disaggregate models will assign a probability of zero to that event and the entire likelihood function will be driven to zero (driving LL to negative infinity).

To test the sensitivity of the parameter estimates to the purchase histories, we use a bootstrap procedure similar to the one described by Chapman and Staelin (1982). We start with the set of parameters estimated from a consumer's actual purchase history and then estimate "pseudovalues" of those parameters by resampling from the actual purchase histories and reestimating the model parameters. Full details of the bootstrapping procedure are given in the Appendix. In essence, we create 50 replicates of each parameter from modified purchase histories, and then examine the *variability* of those estimates and the *differences* between those pseudovalues and the original parameters. In general, if the pseudovalues tend to be tightly clustered and close to the originally estimated parameters, we have more confidence that the original estimates are not particularly sensitive to the types of promotional environments or number of purchases made by each consumer.

Before we discuss the results of the bootstrapping procedure, several notes are in order. First, when analyzing the estimated preference parameters, we choose a random (acceptable) brand for each consumer instead of considering every brand for every consumer. This random set of consumer-brand combinations is used for all analyses involving preference parameters, thereby enabling us to perform univariate analyses instead of the more complicated multivariate tests.

Second, when analyzing the bootstrapped parameters, we replace the DMNL promotion sensitivity parameters, Θ_i, with their reciprocals, $1/\Theta_i$. We do this because the parameter Θ_i can take values from one to infinity. With that wide range of possible values, one would expect great variability in bootstrapped parameter estimates. In contrast, $1/\Theta_i$ takes values from zero to one, similar to γ_i. We believe the across-model parameter stability comparison is more meaningful with the two promotion sensitivity parameters varying over the same range.

Bootstrapping Results

To estimate a parameter's *stability* we calculate the standard error of the mean (SEM) of each parameter's pseudovalues across the 50 bootstrapped samples. We then report the average of these standard errors across the 200 panelists. Our measure of parameter *bias* is the difference between the parameter estimated by using the actual choice history and the average of 50 pseudovalues of that parameter.

We focus on four parameters: promotion sensitivity as estimated by EBA (γ_i) and as estimated by DMNL ($1/\Theta_i$) and brand preference as estimated by EBA (π_i) and as estimated by DMNL (Π_i).

γ_i = panelist i's promotion sensitivity parameter as estimated by using EBA and the panelist's actual choice history

Θ_i = panelist i's promotion sensitivity parameter estimated by using DMNL and the panelist's actual choice history

π_i = panelist i's preference for a randomly chosen brand as estimated by using EBA and the panelist's actual choice history

Π_i = panelist i's preference for a randomly chosen brand as estimated by using DMNL and the panelist's actual choice history

For each of these parameters, we create 50 pseudovalues and then calculate the average and SEM of those pseudovalues. Finally, we create a statistic comparing the parameter estimated by using the actual choice history with the average of its bootstrapped counterparts. Rather than displaying the rules for calculating these averages, standard errors, and differences for each parameter separately, we report the rules once for a generic parameter η. By substituting γ, $1/\Theta$, π, or Π for η, one can infer the rule used to calculate all statistics used in Tables 2 and 3.

$\hat{\eta}_{is}$ = value of parameter η in panelist i's s^{th} bootstrapped sample

$\bar{\hat{\eta}}_i = \sum_{s=1}^{50} \hat{\eta}_{is}/50$ = average value of parameter η in panelist i's bootstrapped samples

$\sigma_{\bar{\hat{\eta}}_i} = \sum_{s=1}^{50} (\bar{\hat{\eta}}_i - \hat{\eta}_{is})^2/(49 * 50)$ = estimate of parameter stability of panelist i

$\Delta\eta_i = \eta_i - \bar{\hat{\eta}}_{is}$ = difference between actual and bootstrapped parameter values for panelist i

Table 2 provides evidence that both individual-level models stand up to the bootstrapping procedure. From the first row of Table 2 we see that the parameter estimates are reasonably stable. Standard errors between .01 and .02 are not exorbitant given an average promotion sensitivity value of about .6 and an average preference value of about .3. Rows 2 and 3 of Table 2 provide further evidence of parameter stability. From the last row of Table 2 we see that, on average, the original parameters are not much higher or lower than their bootstrapped counterparts.

To explore the effect of purchase history length on parameter quality, we consider the correlation of parameter differences (actual value minus average bootstrapped value) with number of purchases in a consumer's choice history. As we would expect, the absolute size of the deviations increases as purchase histories become shorter. This pattern is reflected in the significant negative correlations (α = .01) for all entries in the first row of Table 3. However, these larger differences for short histories are as likely to be positive as they are to be negative, suggesting that shorter purchase histories are not likely to produce biased statistics. We draw this inference from the lack of significant correlation (at α = .05) between purchase history length and the signed value of the actual minus bootstrap difference reported in row 2 of Table 3.

Table 2
PARAMETER QUALITY STATISTICS
(average standard errors of bootstrapped parameter estimates and differences between actual and bootstrapped parameter estimates)

	γ *EBA promotion parameter*	$1/\Theta$ *DMNL promotion parameter*	π *EBA preference parameter*	Π *DMNL preference parameter*
Stability				
Average value of bootstrapped standard error ($\sigma_{\hat{\eta}_i}$) across 200 panelists	.021	.018	.015	.016
Number of panelists (of 200) for whom difference ($\Delta\eta_i$) is within ±.01	96	104	123	124
Number of panelists (of 200) for whom difference ($\Delta\eta_i$) is within ±.10	183	177	195	193
Bias				
Average value of $\Delta\eta_i$ across 200 panelists	−.007	.004	−.005	−.004

In summary, we see evidence that individual-level model parameters become less stable as purchase histories become shorter. However, the purchase histories used in our study seem long enough to produce reasonably stable parameter estimates. Finally, we find no evidence of bias in parameter estimates, nor do we find evidence that a tendency toward bias is related to length of purchase history.

Differences in EBA's and DMNL's Parameters

We see that it is feasible to estimate individual-level models. Both EBA and DMNL perform favorably in comparison with GL. Further, the individual-level models yield reasonably stable, unbiased parameter estimates. More importantly, we find that the task of estimating a noncompensatory model posed by Johnson, Meyer, and Ghose (1989) is a feasible one. In addition, the parameters of the noncompensatory EBA models differ in interesting ways from their compensatory DMNL counterparts. In this section, we highlight those differences.

Table 3
IMPACT OF PURCHASE HISTORY LENGTH
(correlations of purchase history length with parameter differences)

	γ_i	$1/\Theta_i$	π_i	Π_i
Stability				
Correlation of absolute size of difference ($\lvert\Delta\eta_i\rvert$) with number of purchases made by consumer	−.186	−.180	−.290	−.297
Bias				
Correlation of difference ($\Delta\eta_i$) with number of purchases made by consumer	.090	.043	.118	.114

Comparability. One important difference between parameters estimated by EBA and those estimated by DMNL is parameter comparability. Guadagni and Little (1983, p. 207) point out that logit "produces larger utility values in a model that explains more variance than in one that explains less." This observation suggests that analogous logit parameters cannot be compared across separately estimated, individual-level models (unless, of course, the models explain exactly the same amount of variance, a property that does not tend to hold). With EBA, in contrast, the separately estimated individual-level parameters are comparable. Further, because the EBA promotion sensitivity parameter has a probabilistic interpretation, these parameters can be aggregated in meaningful ways. This comparability property is explored in the Implications section.

Identification of stratified preferences. For many panelists, the EBA model assigns extremely small preferences (often on the order of 10^{-10}) for an acceptable brand. Such small preferences indicate that the panelist may have *stratified preferences;* that is, a panelist distinguishes between "preferred" acceptable brands (the set of brands he or she is willing to buy off promotion) and "marginally acceptable" brands (the set of brands he or she will consider only if they are promoted). This notion of marginal acceptability is fully consistent with the EBA choice process. We illustrate this mechanism in Figure 1.

The EBA model suggests that a consumer is "attuned to promotion" with probability = γ. Under the $(1 - \gamma)$ likelihood that the consumer is not screening for promotions, or if there are no promotions available, he or she will restrict choice to the preferred set of brands (situation 1). If the consumer is screening for promotions and at least one brand from the preferred set is on promotion, he or she is in situation 2 and will choose among promoted preferred brands. Only in situation 3 (when the consumer is screening on promotion, no brand in the preferred set is promoted, and one or more brands in the

Figure 1
FLOW CHART REPRESENTATION OF THE EBA MODEL INCORPORATING STRATIFIED PREFERENCES

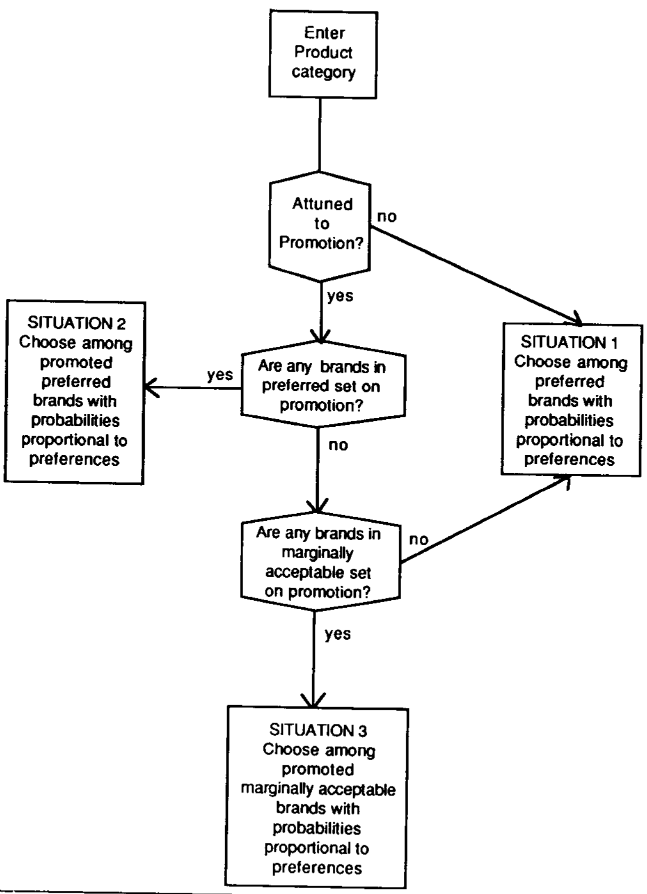

marginally acceptable set is promoted) does he or she choose among marginally acceptable brands.

Because of the noncompensatory nature of the EBA model, a consumer with moderate promotion sensitivity can have these stratified preferences. In contrast, the compensatory DMNL model can show stratified preferences only for consumers who have extremely large promotion sensitivity parameters (Θ's), on the order of 10^{10}, in order to offset the very small preferences assigned to marginally acceptable brands. Such large parameters are possible only for panelists who are "absolutely promotion sensitive" toward preferred brands (i.e., always choosing a preferred brand when at least one is promoted), as is true for 42 panelists in our sample of 200.[8]

[8]In situations when consumers are absolutely promotion sensitive for *all* brands, the EBA model yields γ's of 1 and the two models fit identically in all cases.

Under the DMNL model, consumers with moderate promotion sensitivity can never have stratified preferences. Consider the example we gave before in introducing DMNL. Suppose the consumer has an intermediate level of promotion sensitivity ($\gamma = .6$; $\Theta = 25$, approximately the median values of both parameters in our sample), and suppose that his or her choice history is consistent with brand 1 being in the preferred subset with $\pi_1 = \Pi_1 = 1 - 10^{-10}$ and brand 2 being in the marginally acceptable subset with $\pi_2 = \Pi_2 = 10^{-10}$. The compensatory nature of the DMNL model and its parameters would not allow this stratified preference structure. For example, if brand 2 were ever chosen when on promotion alone, DMNL would be confronted with a term in its likelihood function equal to equation 4. If DMNL assigned the preference values listed above, then

$$D_{2t} = \frac{25 * 10^{-10}}{1 - 25 * 10^{-10}} \approx 0.$$

The maximum likelihood estimation technique cannot allow zero choice probabilities during the estimation period. Thus, in a situation like this (and whenever Θ takes a moderate value), the DMNL model would be unable to separate the Π's by several orders of magnitude and therefore would be unable to identify brands as belonging to a "marginally acceptable" subset.

Notice that EBA equation 6 presents no such problem for a maximum likelihood algorithm:

$$E_{2t} = .6\frac{10^{-10}}{10^{-10}} + (1 - .6)\frac{10^{-10}}{1} \approx .6.$$

The nonlinearities associated with the underlying phased decision strategy enable EBA to assign brand 2 a nontrivial purchase probability despite the fact that the consumer is not absolutely promotion sensitive and the fact that brand 2 is only "marginally acceptable."

Aside from the 42 "absolutely promotion sensitive" panelists in our sample, 46 households appear to have stratified preferences under the EBA model but not under DMNL. We have no way of knowing which (if any) of these panelists actually choose according to a process like that depicted in Figure 1, but the model fit statistics and bootstrap results are somewhat encouraging: of these 46 panelists, 36 fit better under the EBA model than under DMNL, and overall the improvement in fit is significant at the $\alpha = .01$ level. Furthermore, the bootstrapping procedure confirms the stratified preferences 98.7% of the time for these panelists, and for the 112 nonstratified panelists only 17% of the bootstrap runs incorrectly yield stratified preferences.

In summary, we have two individual-level models that both do a good job of predicting choice, despite their different behavioral implications. Both models have stable, unbiased parameters. If we think of DMNL as a linear approximation to the EBA model, it is not surprising that the two exhibit similar fit and predictive capabilities. Despite these similarities, however, the two models' parameters differ in interesting ways, namely in parameter comparability and in the ability to identify stratified preferences. In the next section, we look more closely at these differences, discussing several managerial implications that arise from the EBA model's unique, individual-level parameters.

IMPLICATIONS

Comparability of Parameters

EBA's parameters have the desirable property of being comparable across models. Further, their probabilistic interpretation allows them to be meaningfully aggregated. We now look at one application of this comparability property: differences in promotion sensitivity across brand sizes.

In Table 4 we report an average promotion sensitivity for each brand-size. This value is the average value of γ_i across all panelists who find the brand-size acceptable. From the table, we see that panelists who buy the 3 lb. sizes tend to be more promotion sensitive than those who choose only the 1 lb. sizes. This finding is not unreasonable because larger sizes often offer economic advantage, and the desire to seek economic advantage is probably correlated with promotion sensitivity. The difference in average γ's across these two size groups is significant at the $\alpha = .05$ level.

The only statistically significant pattern within the size groups is that the "specialty brands" (Folgers Flake and Mellow Roast) tend to attract the least promotion sensitive shoppers. Among the major brands, Folgers attracts a slightly, but not significantly, less promotion sensitive following.

This analysis is one example of the benefits of the EBA model's comparability property. A parallel analysis can be performed for other aggregations of the γ's such as differences in promotion sensitivity by store, region, or demographic factors. Similarly, one could develop segments of households based on their promotion sensitivities and examine differences in the brand-preference parameters.

Table 4
AVERAGE PROMOTION SENSITIVITY ACROSS BRAND-SIZES
(average value of γ_i and number of panelists who find indicated brand-size acceptable)

	1 lb. size		*3 lb. size*	
Brand	*Average* γ	n	*Average* γ	n
Folgers	.63	169	.62	94
Maxwell House	.67	138	.71	81
Butternut	.64	129	.71	64
Folgers Flake	.50	32	—	—
Mellow Roast	.50	16	—	—

Stratified Preferences

In this section we use our identification of stratified preferences to explore one way in which promotion affects competition in the marketplace. From a manufacturer's perspective, the existence of stratified preferences suggests that a brand's ability to penetrate the marketplace depends on the brand's promotional status. When unpromoted, the brand will be considered only by consumers for whom it is a preferred brand. When promoted, it will be considered by consumers for whom it is a preferred brand *and* by those for whom it is a marginally acceptable brand.[9]

We look first at the ability of a brand to use promotion to reach additional consumers. We then fine tune that insight by considering the extent to which those additional consumers tend to consider other brands. In some sense, our first illustration shows the extent to which promotion enhances a brand's ability to compete and our fine tuning shows the likely effects of the enhanced competitiveness on other brands.[10]

Figure 2 is a breakdown of panelists' preferences for the eight brand-sizes in our database. The total length of each bar represents the number of panelists who find that brand-size acceptable. The black part of the bar represents the panelists for whom the brand-size is preferred and the shaded part those panelists for whom the brand-size is marginally acceptable.

We see that most of the brands are able to expand their penetration of the panelist pool through promotion. However, some brands expand their penetration more than others. Consider the implications of these data for the manager of Butternut 1 lb. He or she can expand consideration of the brand by more than 40% by promoting.

[9]Note that this statement gives the flavor of consumers' expected behavior in most situations. However, two special cases should be considered. The first case comprises situations in which a preferred brand is not considered for choice (i.e., the probability of choosing the preferred brand equals 0). These situations occur if the consumer is absolutely promotion sensitive (i.e., $\gamma = 1$) and if a marginally acceptable brand is promoted and the preferred brand is not. The second case is made up of situations in which a promoted marginally acceptable brand is not considered for choice. These situations occur if one or more preferred brands are promoted along with the marginally acceptable brand.

Hence the description in the article holds to the extent that consumers are not absolutely promotion sensitive and to the extent that multiple brands do not receive promotional support concurrently. Our analysis of the data suggests that these limiting conditions do occur, but that they are the exception and not the rule.

[10]We could also conduct these analyses on the basis of DMNL parameter estimates. Such analyses would yield less extreme versions of Figures 2 and 3. Recall that DMNL identifies only 42 of the 88 panelists that EBA identifies as having stratified preferences. Because DMNL identifies fewer "marginally acceptable" brands, it necessarily understates the EBA-identified effects for promotion. For example, in Figure 2 we see that Butternut 1 lb. can expand the set of consumers who consider it for purchase by 40% if it promotes. The DMNL analogue suggests that consideration of Butternut 1 lb. could be expanded by less than 20% with promotion.

Figure 2
BREAKDOWN OF PANELISTS' PREFERENCES[a]

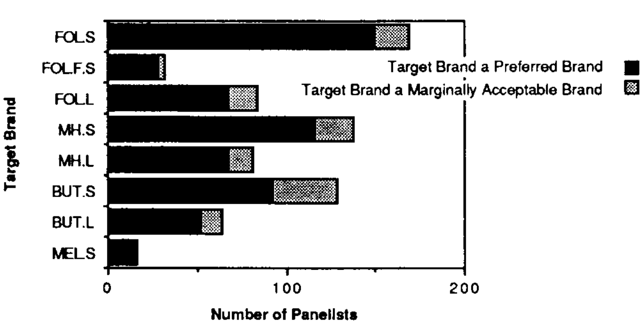

[a]FOL is Folgers, FOLF is Folgers Flake, MH is Maxwell House, and MEL is Mellow Roast. S indicates small (1 lb.) size and L indicates large (3 lb.) size.

No other brand-size in this study can effect such a large gain in consideration through promotion.

In the fine tuning, the manager might explore the other brands considered by consumers who find Butternut marginally acceptable. If those consumers primarily consider other brand-sizes in the Butternut product line, the additional penetration is likely to enhance cannibalization. However, if those consumers primarily consider competitive brand-sizes, the additional penetration is likely to enhance the brand's ability to compete with other product lines.

Figure 3 illustrates the pattern of additional penetration for Butternut 1 lb. In each bar we consider only those panelists who choose the brand associated with that bar. The shadings of the bar represent those panelists' feelings about Butternut 1 lb. (i.e., the first bar suggests that of the 169 panelists for whom Folgers 1 lb. is acceptable, 73 prefer Butternut 1 lb., 41 find Butternut 1 lb. marginally acceptable, and 55 find Butternut 1 lb. unacceptable).

Figure 3
BUT.S STATUS WITH PANELISTS WHO CHOOSE COMPETITIVE BRANDS[a]

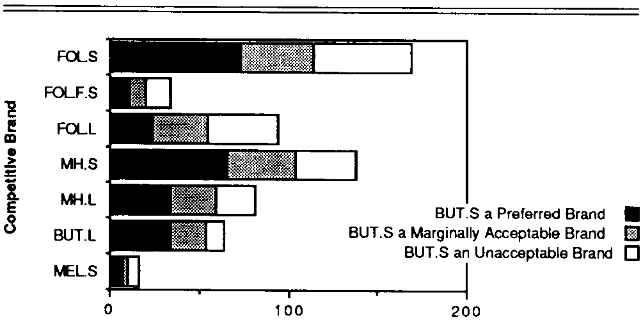

[a]FOL is Folgers, FOLF is Folgers Flake, MH is Maxwell House, and MEL is Mellow Roast. S indicates small (1 lb.) size and L indicates large (3 lb.) size.

Looking at Figure 3 as a whole, we can think of the black parts of the bars as representing Butternut 1 lb.'s penetration of competitive brands' panelist bases in the absence of promotion. The black plus the shaded parts of the bars can be thought of as Butternut 1 lb.'s penetration when Butternut 1 lb. is promoted. We see that promotion increases penetration into Folgers' and Maxwell House's panelist bases more than it increases penetration into Butternut 3 lb.'s panelist base.

Similar fine-tuning analyses could be done for the other seven brand-sizes in the database. From those analyses, Butternut 1 lb. could develop an understanding of the vulnerability of its panelist base to competitive brands' promotions.

In summary, we show that parameters estimated by EBA differ in important ways from those estimated by DMNL. EBA's parameters are comparable across models and its noncompensatory structure allows consumers' stratified preferences to be expressed.

SUMMARY AND CONCLUSIONS

We hypothesize that typical low involvement, repetitive choices in supermarket settings are powerfully influenced by promotional offers. Our EBA model posits a phased decision strategy in which consumers occasionally screen out unpromoted brands. Our approach complements prior noncompensatory models, such as those of Currim, Meyer, and Le (1988) and Gensch (1987).

Rather than using protocol data as often is done in studies of phased decision making, we use UPC scanner panel data to estimate the EBA model's parameters for a collection of consumers. We also estimate parameters for several reference models, including an individual-level compensatory approximation to EBA, termed DMNL. Both EBA and DMNL are shown to predict consumer choices better than the aggregate model of promotion response proposed by Guadagni and Little (1983).

Unsurprisingly, EBA and its compensatory approximation, DMNL, cannot be distinguished in terms of their ability to fit and predict the data. We do show, however, that the comparability of EBA's parameters and the ability of EBA to identify consumers' stratified preferences can provide useful insight.

In sum, we believe the individual-level models derived from EBA are a valuable complement to both aggregate logit analyses and disaggregate, detailed protocol analyses of decision strategies. The disaggregate analysis enables us to capture a phased decision strategy and more completely reflect consumer heterogeneity. In addition, the maximum likelihood procedure used to infer parameters is much easier to implement than protocol analysis. In a matter of seconds we can produce models of hundreds of consumers from their UPC purchase histories. Protocol analysis might take days, or even weeks, to produce a model of a single subject.

Longer purchase histories and richer descriptions of promotional and advertising environments are currently becoming available. With such data and the powerful nonlinear optimization computer software now available, researchers should be able to estimate even more detailed and interesting models of consumer behavior.

APPENDIX
BOOTSTRAP PROCEDURE FOR EXAMINING PARAMETER STABILITY AND BIAS

1. For each of the 200 panelists, estimate their "true" parameters over their entire purchase history (65 weeks). (This step yields parameters γ_i and π_{ik} for the EBA model and Θ_i and Π_{ik} for the DMNL model.) Denote the number of purchases made by each panelist as NUMPRC_i.
2. For each household, create a "simulated" purchase history by the following steps.
   a. Randomly choose (with replacement) NUMPRC_i promotional environments from the set of store environments actually observed by panelist i. This approach allows each purchase history to be altered by repeating or eliminating some observations.
   b. Using the "true" parameters, calculate choice probabilities for all brands by equations 1a and 1b for EBA and equation 2 for DMNL. Then simulate a brand choice decision by using these choice probabilities as weights and choosing a random number.
3. Step 2 provides enough information to rerun the estimation procedure over the simulated purchase history. The results of this reestimation are a bootstrapped γ_i, Θ_i, and two sets of preferences from which we observe π_i and Π_i.
4. Repeat steps 2 and 3 50 times for each panelist. Introduce a new subscript, s, to keep track of the simulation number for each bootstrapped set of parameter estimates: γ_{is}, Θ_{is}, π_{is}, Π_{is}, $s = 1, 2, \ldots, 50$.

REFERENCES

Batsell, Richard R. and John C. Polking (1985), "A New Class of Market Share Models," *Marketing Science*, 4 (Summer), 177–98.

Bettman, James R. (1970), "Information Processing Models of Consumer Behavior," *Journal of Marketing Research*, 7 (August), 370–6.

——— (1979), *An Information Processing Theory of Consumer Choice*. Reading, MA: Addison-Wesley Publishing Company.

——— and Jacob Jacoby (1976), "Patterns of Processing in Consumer Information Acquisition," in *Advances in Consumer Research*, Vol. 3, Beverlee B. Anderson, ed. Ann Arbor, MI: Association for Consumer Research, 315–20.

———, Eric J. Johnson, and John W. Payne (1989), "Consumer Decision Making," in *Handbook of Consumer Theory and Research*, Harold Kassarjian and Tom Robertson, eds. Englewood Cliffs, NJ: Prentice-Hall, Inc.

Chapman, Randall G. and Richard Staelin (1982), "Exploiting Rank Ordered Choice Set Data Within the Stochastic Utility Model," *Journal of Marketing Research*, 19 (August), 288–301.

Currim, Imran S., Robert J. Meyer, and Nhan T. Le (1988), "Disaggregate Tree-Structured Modeling of Consumer Choice," *Journal of Marketing Research*, 25 (August), 153–65.

Davis, Harry L. (1976), "Decision Making Within the Household," *Journal of Consumer Research*, 2 (March), 241–60.

Dawes, Robin M. and Bernard Corrigan (1974), "Linear Models in Decision Making," *Psychological Bulletin*, 81 (March), 95–106.

Gensch, Dennis (1987), "A Two-Stage Disaggregate Attribute Choice Model," *Marketing Science*, 6 (Summer), 223–39.

Guadagni, Peter M. and John D. C. Little (1983), "A Logit Model of Brand Choice Calibrated on Scanner Data," *Marketing Science*, 2 (Summer), 203–38.

Horowitz, Joel L. (1983), "Statistical Comparisons of Non-Nested Probabilistic Discrete Choice Models," *Transportation Science*, 17 (August), 319–50.

Johnson, Eric J. and Robert J. Meyer (1984), "Compensatory Choice Models of Noncompensatory Processes: The Effect of Varying Context," *Journal of Consumer Research*, 11 (June), 528–41.

———, ———, and Sanjoy Ghose (1989), "When Choice Models Fail: Compensatory Models in Negatively Correlated Environments," *Journal of Marketing Research*, 24 (August), 255–70.

Kahn, Barbara E., Donald G. Morrison, and Gordon P. Wright (1986), "Aggregating Individual Purchases to the Household Level," *Marketing Science*, 5 (Summer), 260–8.

Luce, R. Duncan (1959), *Individual Choice Behavior: A Theoretical Analysis*. New York: John Wiley & Sons, Inc.

Lussier, Dennis A. and Richard W. Olshavsky (1979), "Task Complexity and Contingent Processing in Brand Choice," *Journal of Consumer Research*, 6 (September), 154–65.

Manrai, Ajay K. and Prabhakant Sinha (1989), "Elimination-by-Cutoffs," *Marketing Science*, 8 (Spring), 133–52.

Payne, John W. (1976), "Task Complexity and Contingent Processing in Decision Making: An Information Search and Protocol Analysis," *Organizational Behavior and Human Performance*, 16 (August), 366–87.

——— and E. K. Easton Ragsdale (1978), "Verbal Protocols and Direct Observation of Supermarket Shopping Behavior: Some Findings and a Discussion of Methods," in *Advances in Consumer Research*, Vol. 5, H. Keith Hunt, ed. Chicago: Association for Consumer Research, 571–7.

Russo, Jay E. and Barbara A. Dosher (1983), "Strategies for Multiattribute Binary Choice," *Journal of Experimental Psychology: Learning, Memory and Cognition*, 9 (October), 676–96.

Tversky, Amos (1972), "Elimination by Aspects: A Theory of Choice," *Psychological Review*, 79 (July), 281–99.

Reprint No. JMR273105