

Why Modeling Averages Is Not Good Enough

A Critique of the Law of Double Jeopardy

MARTIN BONGERS

Synovate Laboratories
Martin.Bongers@gmail.
com

JAN HOFMEYR

Synovate Laboratories
Jan.Hofmeyr@Synovate.
com

For more than 40 years, Andrew Ehrenberg and others have demonstrated the value of the Dirichlet and related distributions as a method for modeling typical brand performance measures (BPMs) in markets. This work has led to trenchant critiques of many attitudinal measures related to branding—specifically, brand strength, differentiation, and persuasion. Using panel data, we show that though the Dirichlet leads to accurate BPMs for fixed time-slices, its assumptions about individual behavior are wrong. We therefore challenge the way the law-of-double-jeopardy theorists use patterns identified by the Dirichlet to generalize about brand performance, the psychology of consumer preference, and the role of advertising. Far from being problematic, attitudinal measures may be the only path to the successful understanding of the effects of marketing initiatives.

INTRODUCTION

According to the law of double jeopardy (DJ) in markets, big brands benefit doubly: They have more users, and their users use them more. More formally: “In all markets, all conventionally used brand performance measures (BPMs) can be predicted from market share using the Dirichlet distribution (once calibrated)” (Chatfield and Goodhardt, 1975).

The law challenges many marketing-research approaches. Consider brand equity. Most attitudinal systems are based on some sort of commitment or engagement “ladder” (Keller and Lehman, 2004). These are used to identify strong and weak brands—typically by quantifying the percentage of people found at various “levels.” Ehrenberg and others have argued repeatedly, however, that all can be accounted for by brand size alone. To quote “... there are clearly big brands and smaller brands, (but) there is no evidence... that over and above this there are ‘strong’ and ‘weak’ ones” (Ehrenberg and Goodhardt, 2004).

Consider also loyalty programs, or customer relationship and customer experience management. Most commercial approaches focus on

relationships with existing customers as the most efficient way to increase profits (e.g., Reichheld, 1993). Yet DJ implies that retention and share of wallet increase only if penetration also increases. In other words, a narrow focus on current customers will work only if it also leads, whether intentionally or not, to the creation of new customers.

This research does not question the many patterns that scholars have identified using the Dirichlet distribution. The evidence is overwhelming. Yet one must ask why some marketers often behave as if DJ does not exist.

There is at least one reason. As Ehrenberg and Uncles note (1995), the law holds “... for competitive markets which are in a near-steady state”—in other words, so long as market shares do not change too much. Market shares *do* change, however. And so marketers persist in using tools that ignore the law because they hope that by doing so, they will develop fortune-changing brand strategies. The result is an unsatisfactory standoff—between the patterns revealed by the Dirichlet on the one hand and attitudinal brand performance as measured by marketers on the other.

This article examines the empirical veracity of DJ by comparing the statistical assumptions that underpin it with real behavior in panels. We show that the Dirichlet is based on assumptions about individual behavior that are false. In so doing, we show that although DJ accurately describes aggregate market behavior, it does so in a limited sense. Additionally, the authors argue that its empirical flaws can best be resolved by means of attitudinal measures.

A BRIEF INTRODUCTION TO THE DIRICHLET MODEL OF MARKETS

Scope

Commercial operators (Nielsen, IRI, TNS, Gfk, etc.) typically report a wide range of brand performance measures (BPMs)—market share, penetration (i.e., percent buying in a specified time period), the average number of purchases of the brand per buyer, the percent who buy the brand above a certain threshold, and the like (Ehrenberg and Goodhardt, 2004).

Over the years, Ehrenberg and others have shown that all these measures can be predicted by using the Dirichlet distribution with very little market measurement. Implementation takes just two steps: first, a calibration or model estimation step and, second, a calculation step. To estimate a model, you need to specify a “base period” (typically six months or a year) and then make four measures:

- The percent of all consumers (panelists) who buy in the category at least once
- The average number of purchases each category user makes
- The penetration of just one of the bigger brands
- The average number of purchases each of the bigger brand’s buyers makes.

These inputs are used to estimate the model’s parameters for the time period

and market of interest. Once those data are in place, all BPMs for all brands can be estimated with remarkable accuracy just by entering the brand market shares or penetrations.

In terms of reach, Dirichlet patterns have been found across multiple geographies for a multitude of different kinds of market including goods with a very high purchase frequency (e.g., gasoline), durable goods (e.g., PCs and automobiles), and categories in which loyalty programs are thought to play an important role (e.g., airlines, retailers).

Statistical Underpinnings

The model is based on the following assumptions about behavior in markets:

- Each person buys in the category at a steady, long-run rate (e.g., monthly over a six-month period), although a person’s buying rate can fluctuate during the period.
- Although people buy according to a long-run rate, a person’s purchase intervals are irregular (“as if random”) and are independent of the previous interval.
- People are heterogeneous in terms of their category-buying rates. A few buy heavily, but most buy relatively lightly.
- Although brand repertoires and brand-purchase probabilities vary from one person to the next, each person has a steady or fixed likelihood of buying each brand.
- Over time, brands are bought by a person according to that person’s fixed-purchase probabilities. What a person buys, however, seems to be random each time.

There appears to be a good fit between the gamma distribution and the behavior of individual panelists in commercial panels. This leads to the observation (see the

foregoing): “Although people buy according to a long-run rate, a person’s purchase intervals are irregular (‘as if random’) and are independent of the previous interval.” It implies that it is impossible to use previous purchases to judge when a person will buy something again, based on when they bought last.

There also appears to be a good fit between the zero-order Poisson distribution and panelists’ brand buying. This leads to the assumption (see the foregoing), “Over time, brands are bought by a person according to that person’s fixed-purchase probabilities. What a person buys, however, seems to be random each time.” It implies that we cannot tell which brand people will buy next based on what they bought last (although we can know how often they will buy each brand over the specified period). As an example, think of someone who uses each of two brands 50 percent of the time. Their steady purchase probabilities are 0.5 and 0.5. Just as a roulette wheel spins in a random fashion, however, a previous purchase gives no information about which brand someone will buy next.

Psychological Explanations

Ehrenberg and his colleagues have not been shy about making the theoretical leap from the Dirichlet description of market behavior to a consumer psychology that might explain such behavior. They suggest, for example, that most people in mature markets have learned all they are going to about brands. As a result, people learn nothing new from using a particular brand “yet again” or seeing “yet another advert” for it. To quote: “... the experienced consumers’ steady purchase propensities can be thought of as the outcome of years of past experience” (Ehrenberg and Goodhardt, 2004). Well-formed brand opinions—relatively impervious to further marketing—lead to steady purchase

probabilities that can be modeled by zero-order statistical distributions.

The “no-new-learning” theory gains support from the way people behave when brands are discounted. Although price-related promotions can lead to dramatic spikes in brand buying, analysis shows that discounted brands mostly attract previous users. To quote Ehrenberg: Before-to-after repeat buying doesn't change; no before-to-after market share changes occur; and *few* new buyers are attracted. Indeed, the presence of any new buyers is all it takes for the possibility of market-share change. The entire body of evidence suggests that discounting does not induce the kind of use that leads people to change their brand opinions.

In this view, there also is a “no-new-learning” approach to advertising. Like discounting, advertising may cause buying “spikes,” but it works mostly by reminding experienced users of the brand's existence. Mostly, therefore, it attracts people who have used the brand before. It does so by “publicizing” the brand—or “nudging” the brand closer to the front of a person's mental brand queue. By improving the brand's psychological accessibility in this way, it temporarily improves the chances of the brand's being bought. Advertising, from this perspective, does not “persuade” or change minds.

In markets wherein repertoires are possible—and most people do have a repertoire—100-percent loyalty is rare. If loyalty, however, is a matter of degree—for example, varying from 0 to 100—and if people buy the brands in their repertoires according to fixed purchase probabilities, most people are what Ehrenberg calls “typically polygamous.” They are neither “monogamous”—100 percent loyal—nor “promiscuous”—not attached to anything. They buy according to “split loyalties.” This leads to more psychology or, as Ehrenberg and Goodhardt observed

in the *Journal of Business Research* in 2004, “... staying with a repertoire requires less mental effort than making new choices, but having a repertoire enables consumers to exercise some choice without having to reevaluate all brands....”

In spite of the static psychology implicit in the Dirichlet, it is important to understand that DJ does not challenge the existence of marketing levers (e.g., discounting and advertising). What Ehrenberg et al. argue is that people's resulting behavior is “sufficiently idiosyncratic and irregular to be successfully modeled mathematically as being quasi-random, *especially collectively*” (2004, author's italics). In other words, DJ is good at describing aggregate behavior; *that* behavior is consistent with both the Dirichlet assumptions about individual behavior and a DJ-based critique of marketing concepts such as brand strength, differentiation, and persuasion.

Deviations From the Model

Scholars have long noted that market behavior can deviate from DJ—more specifically, that DJ for big brands may be even higher than predicted (Fader and Schmittlein, 1993).

The most important deviation, however, has to do with the model's description of markets as “stationary.” There are two obvious ways in which markets are not stationary. First, market shares may change as a function of the successful introduction of a new brand and, second, market shares may change over time.

Ehrenberg et al. acknowledge that dynamic markets are as yet unexplained (Ehrenberg and Goodhardt, 2004). They argue, however, that the Dirichlet copes rather well with market share changes, because it continues to model fixed time-slices through the change. So, for example, between 1981 and 1992, the market share of Folgers (a coffee brand), doubled in the United States. Dirichlet models fit at both

the beginning and the end of the period (Baldinger, Blair, and Echambadi, 2002).

A CRITIQUE OF THE LAW AND THE IMPLICATIONS DERIVED FROM IT

In reality, the buying behavior of many individual people in markets is dynamic (DuWors and Haines, 1990). This section begins with two examples of such behavior from panel data. The authors show that these are not isolated instances and discuss the implications for the DJ approach to brand-performance estimation and, consequently, for the DJ understanding of how advertising works.

A Look at Two Examples of Real Buying Behavior

In an analysis of a purchase stream of the chocolate purchases of a person participating in a consumer panel (Table 1), the changes in this person's purchase probabilities are clear: This panelist bought chocolate eight times in the first period, nine in the second, six in the third, and five in the fourth. “Brand 1” is the first brand, but Brand 4 dominates for the first 6 months and then almost disappears. Brand 2 gets bought most in the second 6 months but is hardly dominant—and then it, too, disappears. No brand dominates in the second and third periods. Brand 11, however, dominates in the final 6 months—and some might argue that one can discern a developing preference for it in the third period.

For someone such as this, the only way to get “fixed purchase probabilities” (as per DJ and the Dirichlet assumptions) is by *the artifice* of fixing them. So, in period one, Brand 4 has a 50 percent share of wallet (i.e., a 0.50 chance of being bought). Extend the lens to 12 months and both Brands 2 and 4 have a 0.29 chance of being bought. Together, the two account for 59 percent of what this person buys in year one. In year two, *they are not bought at all.*

TABLE 1

A Two-Year Purchase Stream for a Panelist Buying Chocolate*

Purchases	Brands																					
	1	2	3	4	5	6	7	8	9	10	11	12	13									
1	1																					
2		1																				
3			1																			
4				1																		
5					1																	
6						1																
7		1																				
8							1															
9								1														
10			1																			
11									1													
12						1																
13										1												
14			1																			
15											1											
16												1										
17													1									
18												1										
19													1									
20														1								
21															1							
22																1						
23																	1					
24																		1				
25																			1			
26																				1		
27																					1	
28																						1

*With thanks to the Aztec household panel in Australia for supplying all the panelist purchase histories used in this article.

In an analysis of a panelist's purchases of coffee for a year (Table 2), once again brand-purchase probabilities cannot easily be "fixed." In the second six months, the

panelist switches main brand, increases his or her repertoire from three to five brands; and nearly doubles the rate of her or his coffee buying.

TABLE 2

A One-Year Purchase Stream for a Panelist Buying Coffee

Purchases	Brand																
	1	2	3	4	5												
1	1																
2		1															
3			1														
4				1													
5					1												
6						1											
7							1										
8								1									
9									1								
10										1							
11											1						
12												1					
13													1				
14														1			
15															1		
16																1	
17																	1

One cannot look at these panelists' behavior without concluding that their purchase probabilities are dynamic. The only way to make them look "fixed" is by the artifice of taking their average behavior for a specified time. Expanding or contracting the time at will may smooth away the dynamism and "fix" propensities by definition. Such "smoothing," however, does not do justice to the underlying behavior, nor does it do justice to the implicit dynamism of the psychological preferences revealed by the behavior.

From Dynamic Individual Behavior to Non-Stationary Markets

If the behavior of the panelists shown previously were exceptional, there might

not be a problem. For each of four market-leading brands, the authors classified chocolate lists' buying behavior over 24 months and 12 months into one of six mutually exclusive groups as a function of how their purchase probabilities change from one six-month period to another. Using the typical Ehrenbergian definition of a "user" in a period as someone who buys the brand at least once, there are four "user" and two "non-user" classifications:

• **Users**

- *Up*: significantly increase the share of wallet that goes to the brand
- *Static*: share of wallet stays the same
- *Down*: significantly decrease the share of wallet for the brand
- *Defect*: stop using the brand.

• **Nonusers**

- *Adopt*: start using the brand
- *Zero*: remain non-users.

A difference in individual behavior across a time period was defined to be statistically significant if it exceeded the square root of the share of wallet in each of the successive periods.

As an example, consider someone who gives a brand 25-percent share of wallet in the first time period and 36 percent in the second. The confidence interval for the two periods is ± 5 percent and ± 6 percent, respectively. In this case, the difference is not statistically significant, and the person's behavior in relation to that brand is classified as "Static." If, however, the share of wallet in the second period is 49 percent (confidence interval ± 7 percent), the person can be classified as "Up"; Table 3.

By this approach, very few users' brand-purchase probabilities can be described as "near stationary" in either market through any transition. This underlying dynamism need not result in nonstationary markets as long as the changes in each person's

TABLE 3

A Summary of Panelist Behaviors over Time (A: Chocolates, 24 months; B: Coffee, 12 months)

		A			B			A	B
		P ₁ -P ₂	P ₂ -P ₃	P ₃ -P ₄	P ₁ -P ₂	P ₂ -P ₃	P ₃ -P ₄	P ₁ -P ₂	P ₁ -P ₂
Brand 1	Up	138	124	141	31	28	32	20	6
	Static	83	63	60	19	14	14	28	8
	Down	122	145	119	28	33	27	29	8
	Defect	32	48	49	7	11	11	33	10
	Adopt	37	37	50	8	8	11	52	15
	Zero	27	22	20	6	5	5	182	53
	Total	439	439	439	100	100	100	344	100
Brand 2	Up	100	78	100	23	18	23	9	3
	Static	55	56	54	13	13	12	13	4
	Down	83	94	70	19	21	16	14	4
	Defect	59	83	64	13	19	15	35	10
	Adopt	73	60	75	17	14	17	32	9
	Zero	69	68	76	16	15	17	241	70
	Total	439	439	439	100	100	100	344	100
Brand 3	Up	63	59	66	14	13	15	10	3
	Static	42	45	39	10	10	9	11	3
	Down	56	74	68	13	17	15	4	1
	Defect	70	67	71	16	15	16	14	4
	Adopt	84	66	62	19	15	14	21	6
	Zero	124	128	133	28	29	30	284	83
	Total	439	439	439	100	100	100	344	100
Brand 4	Up	8	14	12	2	3	3	0	0
	Static	11	9	10	3	2	2	6	2
	Down	14	14	12	3	3	3	12	3
	Defect	56	29	63	13	7	14	25	7
	Adopt	33	60	45	8	14	10	13	4
	Zero	317	313	297	72	71	68	288	84
	Total	439	439	439	100	100	100	344	100

behavior are offset by changes in others. In fact, market shares in the coffee market were stationary over the period, but the chocolate market was nonstationary.

The result is nonstationary behavior underpinning both a stationary and a nonstationary market. In the chocolate market, the net effect of the behavior is that the market shares of the four leading brands changed. Three of the four brands experienced large gains (from 32 percent to 39 percent; 11 percent to 14 percent; and 10 percent to 12 percent, respectively). The big loser was the brand ranked second at the start of the 18 months: Its share fell from 31 percent to 17 percent—a statistically significant change.

To follow up the original analysis, the authors conducted additional studies in four markets: chocolate, cheese, laundry detergents, and sweets (Table 4). Three of the four were stationary at an aggregate level. The authors focused on the purchase propensities associated with panelists' dominant brands in two successive 12-month periods. Panelists without a dominant brand in either period because of ties were dropped. The reduction was about 14 percent leaving 86 percent of the panelists with clear "winners" in both periods.

The question the authors sought to answer: How many panelists changed their buying behavior enough for that change to be significant? In other words, how many panelists do not have steady purchase propensities? The binomial-distribution method was used to test for significant differences (Agresti and Coull, 1998).

The data showed that many panelists' purchase propensities change. The 24 percent whose brand preference for sweets changed was not as revealing as the much smaller 7 percent who changed their preference for laundry detergents. The low number was not unexpected, given a category marketing history of high levels

TABLE 4

Significant Changes in Purchase Propensities for a Panelist's Dominant Brand

			<i>N</i>	Percent
Cheese <i>(Panelists: 555)</i> <i>(Brands: 163)</i>	Main Brand₁	Significant	60	12
		Not Significant	436	88
	Main Brand₂	Significant	61	12
		Not Significant	432	88
Chocolate <i>(Panelists: 542)</i> <i>(Brands: 53)</i>	Main Brand₁	Significant	66	15
		Not Significant	379	85
	Main Brand₂	Significant	85	19
		Not Significant	363	81
Laundry Detergents <i>(Panelists: 466)</i> <i>(Brands: 55)</i>	Main Brand₁	Significant	26	6
		Not Significant	386	94
	Main Brand₂	Significant	33	8
		Not Significant	379	92
Sweets <i>(Panelists: 499)</i> <i>(Brands: 222)</i>	Main Brand₁	Significant	105	25
		Not Significant	317	75
	Main Brand₂	Significant	103	24
		Not Significant	324	76

of brand commitment (a concept that DJ theorists do not seem to believe in). Seven percent, however, is enough to cause significant market share changes. It is also enough to refute the Dirichlet generalization about individual behavior.

Real Consumers and Dirichlet Fictions

In most stationary and nonstationary markets, there are many people whose purchase probabilities are unsteady. This leads to the question: How can nonstationary individual behavior lead to stationary markets? The answer: It does so as long as the changes in behavior of each person are offset by changes in the behavior of others.

To illustrate the point, the authors studied three consumers buying three brands through three time periods (Table 5a). The three behaved more or less as many of the others had: They changed their buying rates and switched brands. In the first period, the first consumer was a "heavy" buyer, and the other two bought at about half the rate. In the second period, the second consumer became the "heavy" buyer—a change is offset by the first's buying less. In each period, each consumer appeared to change his or her brand preferences, but the changes were offset by changes in the behavior of the others. The three brands maintained market shares of

TABLE 5

Creating Fictitious Dirichlet Consumers (A: Typical Consumers; B: Dirichlet “Fictitious” Consumers, per Our Data)

	A			B		
	1	2	3	D ₁	D ₂	D ₃
Purchases	1			1		
1	1			1		
2	1	3	1	1	3	1
3	2	1		1	2	
4	2	2	1	2	2	1
5	1			1		
6		1	2	1		2
7	2	1		1	2	
8	3	2		2	3	
9	1	1	2	1	1	2
10		1		1		
11	2			2		
12	3	1	1	1	3	1
13	1	2	1	1	1	2
14	1	2		1	2	
15		1		1		

Note: Cells in 5a show which of the three brands was bought by each consumer at each purchase occasion. Cells in 5b rearrange the purchases of the first two consumers to create fictitious consumers whose behavior is “near steady state.” The market is stationary: Brand 1 gets about 60 percent market share in each period, and brands 2 and 3 get about 30 percent and 10 percent, respectively.

about 60 percent, 30 percent, and 10 percent, respectively.

It is a relatively simple matter to create consumers who conform to the Dirichlet assumptions—in other words, consumers whose buying behavior exhibits steady state preferences—by rearranging purchase events (see Table 5). In this way, it is possible to make it look as if a market is populated by consumers whose preferences do not change. And as long as a market is stationary, brand purchases can be assigned to “consumers” in such a way that it looks as if the market is populated by consumers whose behavior is stationary—a situation that may be described as “mutually compensating” changes in

behavior. The reconstruction, however, consists of fictitious characters whom the authors call “Dirichlet fictions.” Additionally, DJ models are underpinned by fictitious consumers.

In Summary

If you fix things by artifact (i.e., by specifying a base time period—six months or a year), individual behavior *can be made* to look static, and market shares can be modeled *as if* stationary. Moreover, as long as all panelists’ changes in purchase probabilities are offset by changes in the purchase probabilities of others, aggregate BPMs based on the assumption of station-

ary individual behavior will not need to be re-estimated as time passes.

In the data analyzed by the authors, however, the statistical assumptions on which DJ is based are inaccurate descriptions of individual purchase probabilities. Nonstationary individual behavior is common enough to cause market share changes. In four of the five markets, individual changes offset one another, and market shares were static. In the chocolate market, however, market shares changed. If the dynamic behavior did not lead to market share changes, one might still argue that it does not matter—that DJ is based on simplified descriptions of individual behavior. BPMs in the market still could be modeled and might have a degree of credibility, but accurate averages produced by a statistical method that does not describe individual behavior, is poor science.

When faced by nonstationary markets, the typical DJ response is to argue that brand performance can be remodeled for time-slices of those markets. Each “slice,” however, would require new parameter estimation; and the resulting BPMs would be particular to each “slice.” The authors believe this is not an efficient approach to brand performance estimation.

THE LANGUAGE OF BRAND PERFORMANCE

In the authors’ research, the statistical underpinnings of DJ provided a poor description of the behavior of many participating panelists. How do these findings reconcile with the psychological explanations that are derived from the idea that individual purchase probabilities are “fixed?” The authors explore three explanations.

A narrow focus on current customers will work only if it also leads, whether intentionally or not, to the creation of new customers.

“Big” and “strong” (or “small” and “weak”) are not the same thing

Over the years, DJ theorists have repeatedly claimed that “there are no strong or weak brands, only big or small ones” (Ehrenberg and Goodhardt, 2004). In dynamic markets, however (e.g., the chocolate market as measured in the authors’ panel), most marketers would want to be able to describe brands not only in terms of their size but in terms of the likely direction of their market share changes.

During the two years covered by the authors’ chocolate data, the number-two brand suffered a sustained fall in market share. If measures had been in place to forewarn of its impending losses, marketers might have described it as “big,” but they almost certainly would not have described it as “strong.” A more likely description would have been “big, but vulnerable.” The same would be true for small brands about to get bigger—they might be described as “small,” but they would not be described as “weak.” In dynamic markets, therefore, big and small aren’t synonyms for strong and weak.

The language of brand strength better describes brands whose market shares are changing. Perhaps one reason the DJ position on the language of brand strength has been relatively easy to defend is that dramatic changes in market share are relatively rare. Perhaps another is that attitudinal BPMs do not have a good record of predicting such changes.

With respect to the second point: Progress is being made when in terms of attitudinal approaches that successfully predict behavioral changes from

point-in-time surveys (Durbach, 2009). With respect to the first, the fact that dramatic changes are relatively rare does not alter the need for a language that appropriately encompasses the possibility of change. If there is a chance that a big brand might be about to suffer a loss in market share, the language of “big but vulnerable” would better describe it than the language of “bigness” alone.

A More Realistic Psychology

Like Ehrenberg, we adopt the idea that a person brings a set of attitudinal brand preferences to each purchase situation. One way to quantify this is in terms of a set of brand-purchase probabilities. There are at least two ways, however, to describe the position of brands in the set.

One uses the idea of “split loyalties”—the idea that people are loyal to each of the brands in the set because they continue to buy each one according to its base probability of being bought. The other is to use the classic language of attitudinal brand attachment to signal differences in the strength of attachment—“strong” when the probability is high and “weak” when it is low. Which you choose does not matter, but we note that the latter is closer to the language most people would use to describe the way the mind works. From this point of view, DJ arguments that challenge the idea of measuring attitudinal brand attachment are mere semantics, nothing more.

The more serious DJ challenge comes from the idea of “no new learning” in mature markets where brands are well known. This idea would be easy to defend

if panel behavior were consistent with “near-steady” purchase probabilities; but it isn’t. Using the method of “revealed preferences”—using actual behavior as a guide to underlying psychological preferences—our records suggest that individual sets of attitudinal purchase probabilities are not stable.

And this leads to the question: What is the cause of the brand switching we see?

Obviously, a great many factors may play a role in the drift of people from one brand to another. A brand may fall out of a person’s repertoire because it is no longer well distributed, because of pack and variant changes, or because of pricing. All such changes involve a consumer’s “learning.” There is one kind of “learning,” however, that DJ theorists argue does not occur: the kind that involves a consumer’s changing his mind about a brand’s characteristics.

On the evidence of the dynamic behavior we show, there is no basis for the claim that “new learning” does not occur. From a psychological point of view, dynamic behavior may involve a person’s changing his or her taste rather than brand opinions. If that were the only explanation, “no new brand learning,” in fact, would have taken place. The same scenario, however, could result in a situation in which a person’s taste had not changed but his or her brand opinions had. Both explanations are plausible, but which one is operative—changing tastes or changing brand opinions—cannot be answered by looking at behavior alone.

One final point: Support for the idea of no new learning comes from the short-run behavior that occurs when people switch brands for a promotion or in response to an advertisement. Some people take longer than this short time frame, however, to drift away from one brand to another—a difference in panel behavior that needs to be analyzed in relation to marketing inputs in greater detail and over

longer time periods. Just because people might have bought a brand in response to an ad does not necessarily mean that this new purchase is (or is not) the beginning of a long-term change in their brand preferences.

Can Individual Purchase Propensities Be Measured by Observation Alone?

In the second year of our analysis, the first panelist in our purchase-stream analysis failed to buy either of the two brands that dominated his repertoire in the first (see Table 1). In the second half of the first year, only on one occasion did he or she purchase the brand that accounted for 50 percent of his purchases in the first half. Such behavior could be expressed as a function of the average propensity to buy each brand during a specified period of time—the “revealed preference” method used as a way to quantify purchase probabilities.

The resulting averages however, would mask what we think is best described as this person’s dynamic drift from one preferred brand to another. In fact, what the behavior “reveals” is that brand preference has changed, and that leads to a question: Is it possible to establish a person’s purchase probabilities at any point in time? If so, is it at all possible to do so by observation?

When a buyer’s behavior is dynamic, the problem with quantifying that individual’s purchase probability is that past purchases no longer represent his or her ongoing attitudes. Using past behavior only produces accurate probabilities when people’s purchase propensities are stationary. And if behavior is mostly dynamic, the method of “revealed preference” is irrevocably flawed.

The point is simple: If dynamic behavior is possible, purchase propensities cannot be established only by observation. That leaves only one alternative: the traditional

Is it possible to establish a person’s purchase probabilities at any point in time?

method of attitudinal surveys. We conjecture that only attitudinal surveys can provide a reliable method to establish how strongly attached a person is to a set of brands at a point in time.

If marketers seek to understand why people have the preferences they do, attitudinal methods are essential because behavioral methods (revealed preference, for instance) cannot overcome the problem of out-of-date individual purchase information (Hofmeyr, Holtzman, Goodall, and Bongers, 2008).

People do have preferences, but those preferences change, and those changes cannot be part of a “steady state.”

IMPLICATIONS FOR ADVERTISERS

The authors believe that consumer psychology built on the back of DJ is inadequate. This premise has at least two important implications for advertisers.

What to Do if You Believe in a World of “Dirichlet Fictions”

If you believe that consumers buy according to “near” steady-state purchase probabilities in stationary markets, your aim should be—as Ehrenberg and his colleagues have often argued—to “publicize” the brand. Brand communications should be about “nudging” the brand to the front of each person’s mental queue. Your focus as an advertiser should be on maximizing both the reach and the frequency of your communication.

It is a world in which you would be “running hard to stay in the same place,” as DJ theorists sometimes say. You would achieve very little, however, for your brand except by accident—or, more specifically, by accidentally “persuading,” even

though you do not believe persuasion is possible.

What to Do if You Believe That Consumers’ Brand Preferences Change

Brand communication is not just about “nudging” already well-known brands to the front of some “mental queue.” Our data show clearly that consumers’ preferences, as expressed in their purchase streams, change. Longitudinal survey data back this up—many people who are reinterviewed using attitudinal metrics will indicate that they no longer favor the same brands as before (Hofmeyr and Rice, 2002).

The power that brand communication and other marketing levers have to shape what people think has been shown with increasing clarity by new measurement tools that allow us to study the brain (Montague et al., 2004). We also know, through the work of neuroscientists that the formation of brand associations involves changes to the brain (Kandel, 2006). Strong brands are built by linking the functional characteristics of the brand through the contexts in which it is used and by the personal goals and values that are important to people.

The message for advertisers and marketers: To truly understand the effect of any campaign, audit the consumers whose behavior *has changed* in response to the campaign. These include previously rare but not new users; true “first-time” users; and defecting heavy users. Understanding the impact of each marketing initiative at the margins—among those who appear to be changing—will lead to understanding of the persuasive effects of that initiative.

When it comes to advertising, the formula for success is relatively simple: Connect the brand to goals and values that matter to people.

CONCLUDING COMMENTS

At the turn of the century, the American Marketing Association asked marketers to choose the best advertising campaigns of the twentieth century. The top five were as follows:

- Volkswagen: “Think Small”
- Marlboro: “Cowboy” series
- Coca-Cola; “The Pause That Refreshes”
- Intel: “Intel Inside”
- Nike: “Just Do It”

Although VW’s “Think Small” campaign shows how an advertising campaign helped to redefine what people wanted, the Marlboro “Cowboy” is a powerful example of a campaign that changed the way people thought about brands. When the Marlboro campaign was devised in 1952, the brand was seen to be a feminine cigarette—“as mild as May” according to the theme line—with only 0.5-percent market share. Leo Burnett put the brand in the hands of cowboys, and the rest is marketing history: Without changing the product, Marlboro became a strong, masculine brand... and the market leader.

One might argue that neither VW or Marlboro is a good example of “new brand learning” because both brands were relatively unknown before the campaign. That would be to forget, however, that the campaigns changed the way people thought about other brands, not just VW and Marlboro. These anecdotes highlight the importance of straightening brand language and fixing the mistaken accounts of consumer psychology that have been proposed by DJ theorists.

The authors’ examination of the purchase streams of panelists in five markets

demonstrate that the assumptions about behavior that underpin DJ models are not universally true. Many individual purchase propensities are not “near stationary.” Does it matter that the assumptions about individual behavior are false as long as the model produces accurate descriptions of brand performance in markets? In the authors’ view, it would matter even if the model accurately described markets over time.

DJ theorists constantly use words such as *near* or *mostly*—as in “behavior is near stationary” or marketing initiatives “mostly attract people who’ve used the brand before.” A close examination of empirical information, however, confirms that behavior is dynamic and that it is the small incremental changes—the material that is excluded from the not-“mostly” in DJ generalizations—that may be at the root of long-term market share changes. They can be initiated by any of the marketing levers available to a marketer, including brand communications.

The authors, therefore, contend that DJ theorists are inaccurate in the numerous conclusions they draw about people, brands, communications, and markets. More specifically, they argue as follows:

- It is an oversimplification to claim that brands are only “big” or “small” and not “strong” or “weak.” Independent of a brand’s size, it needs to be described in terms of the likelihood that its market share might change, and that condition demands a language of brand strength.
- Individual behavior on panels supports the observation that people’s preferences change. Those changes could come from marketing communication. DJ theorists have not made the case for the theory that “no new learning” takes place in mature markets or that advertising does not “persuade.”

- Using behavior to quantify individual brand-purchase probabilities is irrevocably flawed. The only method of measuring a person’s brand-purchase probabilities at a point in time will be found in attitudinal surveys.

DJ, in fact, has utility as a method for describing market aggregates. It loses its efficacy, however, when it is compromised by simplistic consumer psychology and a narrow view of the options available to marketers.

Marketing scientists will surely get better at predicting behavioral and psychological change. In this regard, we offer two further conjectures:

- Methods to *understand* changes in individual behavior and market share will best be found in attitudinal surveys, just as is the case with methods to estimate individual purchase propensities;
- To the extent that covariates can help explain nonstationary markets, those covariates almost certainly will include spending on brand image changing marketing initiatives, including advertising. **JAR**

.....
MARTIN BONGERS is currently a business intelligence specialist for a leading mobile content provider. Martin graduated from the University of Cape Town BSc (Masters) in mathematical statistics while winning the joint prize for the best statistical dissertation on multivariate volatility modeling in South Africa. In 2007, Martin joined Synovate as a statistician and still maintains strong research ties with the Synovate Laboratories.

.....
JAN HOFMEYR (PHD) is currently the international director of innovation at Synovate Laboratories. Before joining Synovate in 2006, he spent 15 years marketing the Conversion Modeltm internationally, a marketing research tool he developed in the 1980s. Since joining

Synovate, Jan has developed numerous practical market measures for Synovate, including the brand value creator (brand equity), connections (copy testing), and true customer view (customer experience). Jan has taught at universities and presented papers and seminars on matters of branding and commitment in most countries of the world.

REFERENCES

- AGRESTI, ALAN, and BRENT COULL. "Approximate is Better than 'Exact' for Interval Estimation of Binomial Proportions." *The American Statistician* 52, (1998): 119–126.
- BALDINGER A. L., E. BLAIR, and R. ECHAMBADI. "Why Brands Grow." *Journal of Advertising Research* 42, 1 (2002): 6–14.
- CHATFIELD C., and G. J. GOODHARDT. "Results Concerning Brand-Choice." *Journal of Marketing Research* 12,1 (1975): 110–113.
- DURBACH, IAN N. "The Use of the SMAA Acceptability Index in Descriptive Decision Analysis." *European Journal of Operational Research* 196, 3 (2009): 1229–1237.
- DUWORS, RICHARD E., and GEORGE H. HAINES. "Event History Analysis Measures of Brand Loyalty." *Journal of Marketing Research* 27, November (1990): 485–493.
- EHRENBERG, A. S. C., and G. J. GOODHARDT. "Understanding Brand Performance Measures: Using Dirichlet Benchmarks." *Journal of Business Research* 57, 12 (2004): 1307–1325.
- EHRENBERG, A. S. C., and M. D. UNCLES. "Dirichlet-type Markets: A Review. Part 1: Patterns and Theory." Working paper, South Bank Business School and Bradford Management Centre, 1995.
- FADER, P. S., and D. C. SCHMITTLEIN. "Excess Behavior AI Loyalty for High-Share Brands: Deviating from the Dirichlet Model For Repeat Purchase." *Journal of Marketing Research* 30, November (1993): 478–493.
- HOFMEYR, J. H., and J. (BUTCH) RICE. *Commitment-Led Marketing*, Chichester, U.K.: Wiley and Sons, 2000.
- HOFMEYR, J. H., P. HOLTZMAN, V. GOODALL, and M. BONGERS. "A New Measure of Brand Attitudinal Equity Based on the Zipf Distribution." *The International Journal of Marketing Research* 50, 2 (2008): 181–202.
- KAHN, B. E., and D. G. MORRISON. "A Note on 'Random' Purchasing: Additional Insights from Dunn, Reader and Wrigley." *Applied Statistician* 38, 1 (1989): 111–114.
- KANDEL, E. *In Search of Memory: The Emergence of a New Science of Mind*. New York: W.W. Norton and Co., 2006.
- KELLER, KEVIN L., and DONALD R. LEHMANN. *Brands and Branding: Research Findings and Future Priorities*. Special Report 05-200b. Cambridge, MA: Marketing Sciences Institute, 2004.
- MCCLURE, S. M., J. LI, D. TOMLIN, K. S. CYPERT, L. M. MONTAGUE, and R. MONTAGUE. "Neural Correlates of Behavioral Preference for Culturally Familiar Drinks." *Neuron* 44, October (2004): 379–387.
- MONTAGUE, R. *Why Choose This Book?: How We Make Decisions*. New York: Penguin Group (USA) Inc., 2006.
- REICHHELD, FREDERICK F. "Loyalty-Based Management." *Harvard Business Review* March-April (1993): 64–73.