BREAKOUT BRANDS

WHY SOME BRANDS TAKE OFF... AND OTHERS DON'T

JARED SCHRIEBER

BREAKOUT BRANDS
Why Some Brands Take Off...and Others Don't

ISBN 978-1-5445-3505-0 *Hardcover*
978-1-5445-3504-3 *Ebook*

CONTENTS

APPENDIX

THE SCIENCE BEHIND CONSECUTIVE REPEAT RATES AND MARKET SHARE

Coauthored by Joel Rubinson *and* Jared Schrieber

The primary finding of this study, and the core of the Brand Growth Flywheel, is the central role of the Repeat Loop—that is, the frequency with which consumers repeat-purchase a brand dictates the brand's growth or decline. That repeat-purchase frequency improves when consumers become more likely to repeat-buy the brand on their very next occasion for that type of product (versus switching away for a few occasions and then coming back). The percentage of time that consumers repeat-buy the same brand on the very next occasion is called the consecutive repeat rate, or, among us math nerds, the Markov Repeat Rate. But what causes the consecutive repeat rate for a brand to increase?

Consecutive repeat rates increase when consumers' underlying probability of purchasing the brand increases. Let's start with the simple example of a single consumer (you) and a single brand (Ruby's BBQ sauce). Suppose you have a 20 percent probability of choosing Ruby's when you buy barbecue sauce. In that case, you have roughly

a 4 percent probability (20 percent times 20 percent = 4 percent) of choosing Ruby's on two consecutive shopping occasions.[1] However, if you have a 50 percent probability of choosing Ruby's BBQ sauce, then the odds of two consecutive purchases being Ruby's will jump from 4 percent to 25 percent (50 percent times 50 percent). To capture 50 percent of your barbecue sauce purchases, Ruby's will have to get you to consecutively repeat 25 percent of the time. This is why brands with higher market shares have consecutive repeat rates significantly higher than their competitors.[2] And while we can't yet plug a Neuralink into consumers' brains to figure out their purchase probability for a given brand, modern consumer purchase panels allow us to measure their consecutive repeat rate as a mathematical proxy. Moreover, we can track this behavior across consumers and brands to understand what is driving changes in market share within a product category. This starts with a simple brand switching matrix (a.k.a. Markov Matrix) common to "Gained–Lost–Retained" analyses. Let's start with a very simple example where the barbecue sauce category is limited to just three competing brands with stable market shares.

In Figure A1, Morty's had 60,000 buyers on the first purchase occasion, of which 45,000 purchased Morty's again on the very next occasion (75 percent consecutive repeat rate). In addition, Morty's won 11,250 purchases from consumers who had purchased Ruby's on their first occasion and another 3,750 consumers who switched from

1 This oversimplified example assumes no influence of purchase probability from one occasion to the next. While we do not believe this to be entirely true, it remains a hotly debated topic among leading marketing scientists. Regardless, this "zero-order" assumption does not affect the mathematical certainty that brands with a higher probability of purchase must have a correspondingly higher consecutive repeat rate. If, in fact, a purchase occasion does influence the shopper's choice on the very next occasion, then the role of consecutive repeat rate in explaining market share changes is even greater than what we have outlined in these simple examples.

2 Consecutive repeat rate for a given brand is defined as the percent of consumers who repeat buy the exact same brand on their very next purchase occasion for the category.

Reggie's. This pattern is typical in categories where brand switching is balanced and market shares remain stable.

		Second purchase occasion				
		Morty's BBQ sauce	Ruby's BBQ sauce	Reggie's BBQ sauce	Totals for first occasion	Market Shares for first occasion
First purchase occasion	Morty's BBQ sauce	45,000	11,250	3,750	60,000	60%
	Ruby's BBQ sauce	11,250	16,880	1,870	30,000	30%
	Reggie's BBQ sauce	3,750	1,870	4,380	10,000	10%
	Totals for second occasion	60,000	30,000	10,000	**100,000**	
	Market Shares for second occasion	60%	30%	10%		100%

Figure A1. Brand Switching Matrix with Steady Market Shares

What if Morty's suddenly lost 8,000 MORE buyers to Ruby's? Its consecutive repeat rate would decline to 62 percent (37,000 repeats out of 60,000 opportunities to repeat) and its market share would decline to 52 percent (52,000 Morty's purchases out of 100,000 category purchases). This would indicate that the market is no longer in equilibrium and that market shares for each brand are moving toward a new balance of power.[3]

So, when a brand's consecutive repeat rate changes, the brand's market share will also be changing. This is an arithmetic certainty that will lead us to powerful conclusions about the nature of the underlying force that causes both metrics to move together. The bigger the brand's market share, the higher its consecutive repeat rate must be! We see this in real-world data across all product categories, just like in this example.

3 This new stabilized set of shares can be estimated via linear algebra using eigenvectors.

Figure A2. Market Shares and Consecutive Repeat Rates Derived from Figure A1

	Market Share	Consecutive Repeat Rate
Morty's	60%	75%
Ruby's	30%	56%
Reggie's	10%	44%

BREAKING IT DOWN

The figures above represent aggregated data patterns, but what is the underlying process that generates this relationship between market share and consecutive repeat rate? By understanding the micro level, we will garner greater insights and intuitions about how brands win (and lose) market share at the macro level.

Let's bring back our very first hypothetical example of Ruby's BBQ sauce, where you have a 20 percent probability of purchasing the brand, and every other consumer has their own probability of purchase. Is there a process that explains how each consumer can have their own probability of purchase, yet the aggregate of those probabilities results in the kind of market shares and consecutive repeat rates we see in the real world?

There must be an underlying distribution of consumer purchase probabilities for each brand, going from 0 percent to 100 percent. The shape of each brand's probability distribution curve explains the market share and consecutive repeat rates measured for each brand in the category.

Fortunately, there is one type of probability distribution that best explains what we see in real-world data. It is called the beta distribution and it is widely used throughout all the sciences as a probability distribution of probabilities. In our industry, it demonstrates what percent of category buyers have close to a 0 percent chance of buying a given brand, versus a 30 percent chance, versus near 100 percent loyalty.

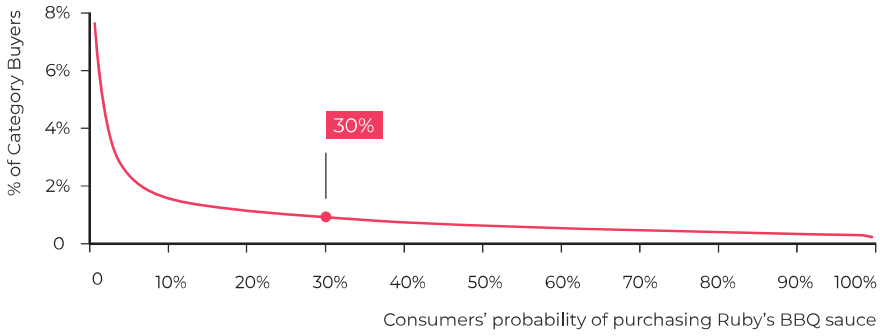

Figure A3. Beta Distribution for Ruby's BBQ Sauce at 30 Percent Market Share

The beta distribution has two parameters, alpha and beta, which have very natural interpretations for marketers. A brand's market share is the same as the mean (average) value of its beta distribution curve, which is a function of the ratio of alpha to beta (specifically by dividing alpha by the sum of alpha + beta). In our barbecue sauce example, we know that this ratio must be 30 percent to equal Ruby's market share. However, there are infinite possibilities for alpha and beta that would give the same ratio (e.g., .3/.7; 3/7; 6/14). How can we choose? That is where the consecutive repeat rate comes in.

In statistics, there are two things we always want to know about the shape of a probability curve: its mean and its variance. We established that the beta distribution's mean is a function of the brand's market share. Similarly, the beta distribution's variance is a function of the brand's consecutive repeat rate. Here is how we know this. The variance of any probability distribution is $E(x^2)- [E(x)]^2$. When applying the beta probability density function to consumers' probabilities of buying a given brand, these terms take on intuitive meaning.

The expected share of buyers is E(x), and the expected consecutive repeat rate is just as formulaic; it is $E(x^2)/E(x)$. So, if you know a brand's market share and its consecutive repeat rate, you can solve for alpha and beta and get the whole curve! That's how powerful knowing the consecutive repeat rate is!

Fortunately, you don't need to memorize these formulas to benefit from them. What you do need to know is that a brand's beta distribution curve (which quantifies the probabilities of consumers buying

the brand) can be derived from the brand's market share and consecutive repeat rate. This is pretty remarkable when you think about it; knowing two statistics about a brand allows you to predict with 99 percent accuracy the percent of consumers who have a specific probability (e.g., 80 percent odds) of buying your brand![4] That's not because market share and consecutive repeat rate dictate the beta distribution, but because the probabilities of consumers buying a given brand drive all three. Put another way, a brand's market share cannot move independently from consumers' underlying purchase probabilities or their consecutive repeat rates.

A brand's market share cannot move independently from consumers' underlying purchase probabilities or their consecutive repeat rates.

Figure A4. Alpha and Beta Can Be Derived from Market Share and Consecutive Repeat Rate

	OBSERVABLE MEASURES		DERIVABLE ATTRIBUTES	
	Market Share	Consecutive Repeat Rate	alpha	beta
Morty's	60%	75%	.99	.66
Ruby's	30%	56%	.5	1.16
Reggie's	10%	44%	.17	1.49

4 The authors modeled the beta curve based upon actual Numerator consumer purchase data for forty-five brands, across laundry detergents, nutrition bars, toothpaste, and frozen pizza. The correlation of the beta curve predictions to the observed share of wallet across nine hundred observations was literally 99 percent (using 5-percentile bands for each prediction bucket; 0–5 percent probability of buying, 5–10 percent, etc.).

In fact, it's not just market share and consecutive repeat rate that emerge from this underlying distribution of consumers' purchase probabilities. Brand purchase cycles, any type of repeat rate measure, and household penetration are all emergent outcomes of consumers' purchase probabilities as demonstrated by the beta distribution.

BREAKING OUT

What happens when market shares are not stable and a brand begins to grow? We have already seen that it implies a switching imbalance and that consecutive repeat rates must be growing for the winning brand and declining for any losing brands. Now let's apply our newfound knowledge to dive one level deeper. If a brand's market share and consecutive repeat rate are increasing, then the vector of purchase probabilities across all consumers of the category must have "shifted to the right"—a higher proportion of shoppers are now more likely to buy the brand.

As shown in Figure A5, for Ruby's to increase its market share from 30 percent to 50 percent, the beta curve representing consumers' probabilities of purchasing Ruby's must also shift. The most noticeable shifts occur at the tails of the curve. On the left side of the chart, the proportion of category buyers with low odds of buying Ruby's shows a drastic reduction. This will play itself out in terms of increased buyer penetration. On the right side of the chart, we see a large increase in the proportion of category buyers with a high probability of purchasing Ruby's. This equates to the brand having significantly more highly loyal consumers. This is the scientific explanation for why market share increases always come from having both greater penetration and greater loyalty.

When the beta distribution shifts (as we just saw in the case of Ruby's), the brand's consecutive repeat rate must simultaneously change correspondingly. While we cannot directly observe the distribution of consumer purchase probabilities, we *can* observe changes

in a brand's market share and its consecutive repeat rate. When those measures change, the probability curve must have shifted because those are the only two inputs required to solve for the beta distribution's parameters.

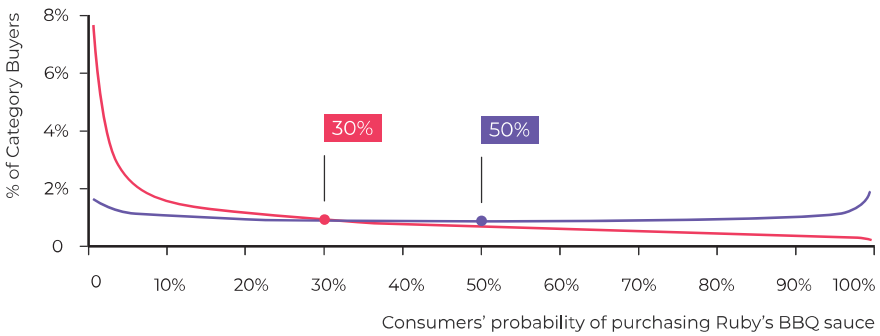

Figure A5. Ruby's Beta Distribution Curve Shifts as It Gains Share from 30 Percent to 50 Percent

If a brand's curve is shifting, and all measures (including purchase cycles, household penetration, and loyalty) change as a result, why should we prefer one metric as a stronger indicator than any other? Which metric is the best indicator for separating a temporary perturbation or measurement failure from a trend?

Consecutive repeat rate offers additional benefits versus the other measures when it comes to distinguishing short-term fluctuations in a brand's share from longer-term trends. This is because traditional metrics such as market share and household penetration are measures that occur in "calendar time." This means that a market share gain in one month could reflect short-term marketing forces (such as advertising campaigns and temporary price reductions) that might vanish or reverse the next month. On the other hand, the consecutive repeat rate occurs in "event time," that is, it reflects a smoothing over weeks or months (or even years) of consumers' probability of buying the brand twice in a row. It is more likely to reflect fundamental, nonfleeting shifts in the underlying distribution of probabilities of consumers purchasing the brand.

PENETRATION IS *NOT* PREDICTIVE

There is a widespread belief in our industry that the most predictive path to market share growth is to grow your brand's buyer base, or "household penetration" (i.e., the percent of consumer households buying your brand at least once in a year or some other specified time frame). The problem with penetration as a predictive measure is that it is not fundamental to the generating process; it is an emergent outcome of consumers' purchase probabilities. When it is hot and sunny, more people buy ice cream; a rise in ice cream sales is, therefore, an emergent outcome of the weather. Buying ice cream does not cause hot and sunny weather! However, hot and sunny weather does increase consumers' probability of buying ice cream, which in turn increases penetration rates for all ice cream brands. Consecutive repeat rates, however, are immune to the weather and other seasonal forces. Furthermore, if a brand's market share, consecutive repeat rate, and the overall category's purchase cycle are known, then the brand's penetration can be accurately estimated within 0.2 percent.[5]

Penetration is an emergent property of consumers' underlying purchase probabilities. Without a proper understanding of this generating process, it's too easy to mix up cause and effect.

BETA DISTRIBUTION VERSUS THE DIRICHLET DISTRIBUTION

The Dirichlet distribution used by many marketing scientists is often characterized as a multivariate beta distribution. In other words, the beta distribution models one brand at a time while the Dirichlet models all brands within a category simultaneously. However, we find

5 Mobile Marketing Association, "Outcome-Based Marketing v2.0." This R&D effort, led by the Mobile Marketing Association, documented the accuracy of estimating penetration by knowing the underlying generating process.

that the Dirichlet is overly restrictive in its assumptions, which hinders marketplace understanding.

The Dirichlet model has a hardwired assumption that all brands in a category are in direct competition. In reality, that is rarely the case. Instead, most well-developed markets are characterized by subsets of brands being more directly competitive within the subset than across subsets—just as we saw in the switching data (e.g., Interaction Index) that measured how winning brands won more than their fair share from smaller, niche brands. For example, Sensodyne toothpaste competes more directly with other sensitive-gum toothpastes than with all other toothpastes. Secondly, the Dirichlet distribution forces us into modeling based on the assumption of market equilibrium. Since we are also interested in spotting departures from equilibrium, fitting individual beta distributions for each brand offers a more agile way of spotting when a brand is breaking out (or breaking bad).

KEY TAKEAWAYS

- **Know what equilibrium looks like.** To identify when a brand is breaking out, we must first establish the baseline of what equilibrium looks like regarding market share and consecutive repeat rate.

- **Market share gains can only be sustained by brands with corresponding increases in their consecutive repeat rate.** Temporary increases in penetration will not sustain market share growth if repeat purchases do not follow. Increases in consecutive repeat rates can be used to predict new, sustainable market share levels.

- **Consecutive repeat rate is the best early signal of brand growth.** Of the three primary measures of brand performance, market share, household penetration, and consecutive repeat

rate, the last one is the best leading indicator of a trend in brand growth. That is because market share reporting is a time-based metric that is retrospective and susceptible to noisy fluctuations caused by temporary marketing programs such as advertising campaigns and trade promotions. Penetration can change for reasons other than gaining or losing market share (e.g., category seasonality). Consecutive repeat rate is native to the generating process and is calculated whenever a consumer makes their next purchase of the category, making it a longer-term measure that reflects a blend of marketing conditions unbound by time.