OUTCOME-BASED MARKETING v2.0

Profitable Growth by Targeting Consumers in the Movable Middle

MMA | neustar

Table of Contents

Executive Summary

When CMOs review the success of their advertising strategy, they measure it against a specific business outcome. Sometimes, that business outcome relates to brand building, and they use metrics like awareness, consideration, or preference. Other times, it relates to customer acquisition or sales performance, and they use metrics like ecommerce checkouts, in-store purchases, loyalty program sign-ups, or app downloads. Yet despite that wide variety in potential business outcomes, many marketers still structure their media plans to maximize reach for the dollars on hand.

Reach, however, is a means to an end. In today's economic climate, marketers are under pressure to deliver media plans that are better aligned with business priorities (e.g., sales lift, coupon downloads, phone calls, foot traffic). That's where outcome-based planning comes in. In outcome-based planning, channel allocation and targeting decisions are designed to help achieve business outcomes, and reach on its own isn't a top priority.

In this joint project between MMA Global and Neustar, we used return on ad spend (ROAS), a widely used measure of marketing success, to develop an innovative 'v2.O' framework—as this report's title notes—to help brands target a new group of consumers: the 'movable middles.'

The movable middles are open-minded consumers defined by their mid-range probability of buying the brand. They're not people who purchase the brand's products at every turn, but they don't have anything against the brand either. They're different for every brand, and on the spectrum between never-buying to always-buying the brand, they're in the middle: yes, they may buy the brand some of the time (and even buy it a lot), but many of their purchases go to the brand's competitors as well.

We've developed a new methodology to identify and target these most valuable consumers and put it to the test for a brand of frozen pizza in the U.S. market. We found that the movable middles for that brand were 5X more responsive to the brand's advertising. We developed an outcome-based marketing plan to target that group, and it outperformed a traditional reach plan by more than 50% on ROAS. It even improved reach across the board.

We believe that this work paves the way for an exciting new framework in marketing strategy.

MMA | neustar

Key Insights

1   **The movable middles are a key new target group.** By trimming off consumers who are very unlikely to buy from a brand as well as those who are very likely to buy from it (and therefore can do with less advertising), marketers can zero in on the consumers that are most responsive to the brand's advertising. For the brand in our study, ROAS for the movable middles was 5X greater than for the rest of the population.

2   **The movable middles are a substantially larger group than heavy buyers,** giving marketers interested in multiple outcomes—like ROAS and market penetration—a much bigger target to work with.

3   **The movable middles can be identified via modeling and targeted via segmentation.**

4   An outcome-based plan that is optimized to target a brand's movable middles is shown to **outperform a traditional reach plan by over 50% on ROAS while achieving higher penetration.**

5   **Outcome-based marketing plans can reach receptive new buyers** and pave the way for future brand growth.

Introduction

In marketing, as in life, simple questions don't always have simple answers.

Since the dawn of marketing, brands have spent considerable time and resources to understand what drives consumers to purchase their products. Modern academic research in consumer behavior and decision theory dates back to the 1950s and 60s and led to real breakthroughs early on, but there was no scanner data at the time to put many of those theories to the test. Even when consumer purchase data came to life in the 1970s and 80s, brands had no way to tie that data to consumers at scale and use it to target them one-to-one. Mass-market channels like TV and radio dominated the media landscape.

Big data and addressable channels changed everything. For the first time, marketers could pinpoint what particular consumers they wanted to reach, up to the individual. But they quickly wrestled with a deceivingly simple question: Is it more productive to acquire new customers, or to get more out of loyal customers?

The answer is neither. In this report, we're taking a novel approach based on propensity analysis, and focusing on a receptive new target audience: the movable middles. They're defined as consumers who are active in the product category (we know they buy frozen pizza, for instance) and have a 20-80% propensity of buying the brand's product.

You might be familiar with the terms 'persuadables' or 'brand switchers' from past literature in this field. Although similar in concept to the movable middles, those other groups are defined squarely by past buying behavior: switchers are medium buyers of the brand, and persuadables are heavy buyers who, based on buying frequency, are getting ready to buy again. **Movable middles are different because they cut across all levels of past purchase behavior. Some may never have bought the brand before.**

How do we quantify their purchase propensity? Our work doesn't rely on massive surveys or harvesting of social data, but rather on combining well-established probability models for customer-base and ad-response analysis with best-in-class audience segmentation. We found that marketing plans organized around the movable middles yield better results than standard plans optimized for reach:

- **More than 50% lift in return on ad spend (ROAS)**
- **Better market penetration across all buyer types**

To understand how we defined and identified movable middles and achieved this level of campaign performance, let's first go over some important market and methodological background.

Marketing plans organized around the movable middles yield better results than standard plans optimized for reach.

New vs. Loyal Consumers

Since we're proposing a model that takes the spotlight away from thinking of consumers as either new or loyal customers, it's helpful to spend a moment to go over the reasons why those two groups developed such a strong following over the years.

Loyal buyers already have a connection to the brand. They have it on speed dial and need very little prodding to purchase it again. Their repeat business accounts for a disproportionate amount of sales, so they're an opportunity for marketers to spend comparatively little (less so than to convince a new customer to try the brand in the first place), maintain the status quo, and get a solid return on that marketing investment.

But too often, loyal buyers are defined as heavy buyers of the brand, or even a subset of those heavy buyers. That's generally a very small group of consumers, even for large established brands, and they invariably turn over—drawn away by competitors, macroeconomic conditions (like today's pandemic), or changes in life stage. Young parents, for example, eventually stop purchasing diapers.

There is a school of thought that marketers who put all of their energy into current buyers and ignore potential new buyers are missing a big opportunity to grow their brand, especially if those new buyers are otherwise regular shoppers of the product category. Market penetration, they argue, is the antidote to double jeopardy.[1]

The academic debate between new and loyal customers can be polarizing. Two camps have formed: one advocating to optimize media plans to target heavy buyers at the expense of every other consumer segment; and the other dismissing targeting altogether and pushing for media plans that maximize reach across the entire population—or at least the population likely to be active in the brand's category, whether it's coffee, frozen pizza, auto insurance, or fast food.

In practice, this has left brands scrambling between two very different approaches to media planning. We believe that's a false dichotomy, and the work presented in this paper offers a new framework focused instead on *responsive* audiences.

---

1 Double jeopardy refers to the idea that small brands in mature markets are at a double disadvantage: they have fewer customers to begin with, and those customers are also generally less loyal too. The concept was introduced by Andrew Ehrenberg fifty years ago, and has found renewed interest through the work of Byron Sharp recently.

Targeting Today

As you will see in this report, our Outcome-Based Marketing 2.0 framework makes use of audience targeting. But isn't targeting controversial today?

Individual privacy has become paramount in consumer markets, and regulatory efforts like Europe's GDPR and California's CCPA are paving the way for stronger consumer protections around the world. As a result, technology companies are doing away with third-party cookies, mobile ad IDs, and other technical workarounds like fingerprinting[2] that have been the backbone of digital targeting for the past twenty years.

This doesn't mean, however, that measurement and attribution are now a lost cause, and that marketers should automatically retreat to broad reach campaigns. In fact, in today's economic climate, a clear brand strategy coupled with effective market analytics are the top-two most vital marketing capabilities for marketers to nurture.[3] The current changes are a much-needed reset for the marketing ecosystem, and an invitation to evolve to a more robust and trustworthy environment.[4]

Purchase and media usage data will still be collected in the future, with people's expressed consent, and marketers are already gaining access to better systems to validate their first-party data and enrich it with reliable second-party data. That data has tremendous value in the moment, but thanks to recent advances in the field of predictive analytics, it's also increasingly valuable to anticipate future behavior.

Done correctly, audience targeting has a bright future.

Done correctly, audience targeting has a bright future.

2 Fingerprinting is a technique whereby a combination of IP, user-agent and JavaScript attributes are used to identify unique users with a high level of accuracy. It's been around for a while as a backup system for cookies, but all the top browser companies are now cracking down on it for the same reasons they're cracking down on cookies and mobile ad IDs.
3 Gartner: The Annual CMO Spend Survey, Part 1
4 Forbes: What digital marketing will look like once the cookie is dead

MMA | neustar

Models Used for Our Research

Purchase Probability Model

There's a vast body of literature on customer-based purchase probability modeling, and we invite you to read the Appendix section at the end of this report for a brief introduction to the key models that paved the way for our approach.

For this study, we focused on consumers' propensity to purchase the brand, and we adopted the beta distribution[5] to model that purchase propensity among category buyers. We selected that distribution over other candidates because of its intuitive simplicity, its historical success in modeling such behavior, and most importantly its near perfect correlation (99%) with the distribution of real-life purchase data in our tests (see the section on Key Processes below).

As a model of brand loyalty, the beta distribution makes the assumption that consumers' purchases within a product category are independent: for example, if someone has a 50% probability of buying the brand, then they have a 25% (50% x 50%) chance of buying it twice in a row. That independence can be debated, but it's largely verified in stable consumer markets, and the beta distribution was an excellent fit for our real-life data.

Using the beta distribution, we can deterministically calculate the distribution of consumer probabilities for the brand using only the brand's market share and its repeat-purchase rate (the percent of consumers buying the brand twice in a row). The beta curve can mathematically take many different shapes, but in consumer research it's generally a U-shape or inverted J-shape, as shown in figure 1a.

For a small brand (like the brand in figure 1a), it shows that a high concentration of category buyers never buy the brand, and that very few consumers are guaranteed buyers. For brands with large market shares (like the brand in figure 1b), it shows that there are also many category buyers who are highly likely to buy the brand. Those observations make intuitive sense and make the beta distribution a good representation of the 'geometry' of the brand.

Figure 1c shows that the beta curve is an excellent match for the real-life distribution of purchase probability for a brand of frozen pizza with a 10.3% market share.
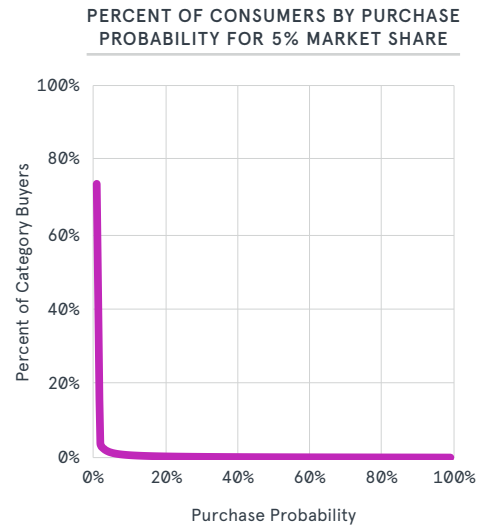
5 Stewart "The Beta Distribution as a Model of Behavior in Consumer Goods Markets"
  Management Science, Vol. 25, No. 9, 1979, pp. 813–821
6 Source: Numerator

PERCENT OF CONSUMERS BY PURCHASE PROBABILITY FOR 5% MARKET SHARE

FIGURE 1A: Beta representation of the distribution of purchase probabilities for brands of different size

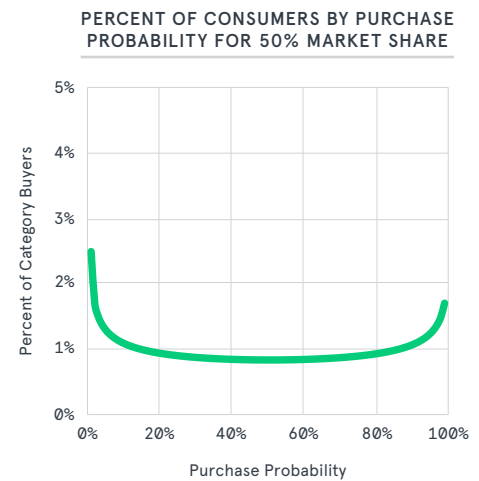**PERCENT OF CONSUMERS BY PURCHASE PROBABILITY FOR 50% MARKET SHARE**

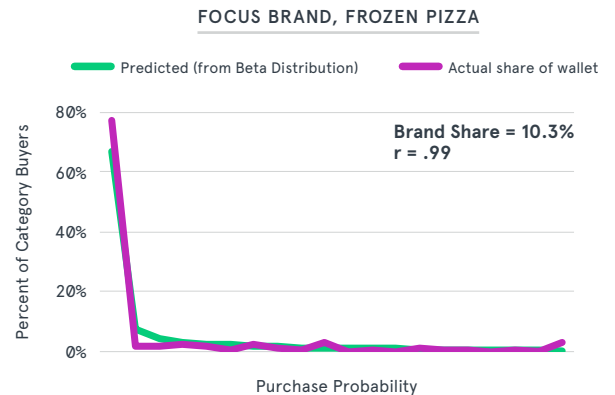**FIGURE 1B:** Beta representation of the distribution of purchase probabilities for brands of different size

FOCUS BRAND, FROZEN PIZZA

Predicted (from Beta Distribution) — Actual share of wallet

Brand Share = 10.3%
r = .99

FIGURE 1C: Beta and real-life distribution for a brand of frozen pizza with a 10.3% market share.[6]

MMA | neustar

Ad Responsiveness Model

A key breakthrough of this project is that we used purchase probability to mathematically connect market share, repeat rate, and ad responsiveness.

Substantial evidence has led researchers over the years to use logistic regressions to model a consumer's responsiveness to advertising. The logistic regression is the workhorse of modern multi-touch attribution models, and it results in an S-shaped utility response curve between marketing exposure and the probability of buying the brand.

While marketing campaigns undeniably have an effect on outcomes, the most predictive variable of future purchase is prior purchase of the brand,[7] and the S-curve does a good job of illustrating that inertia: people who rarely buy the brand are unlikely to start buying it a lot, so advertising will have minimal impact; and people who already buy it often are likely to continue to do so even without much additional advertising.

The consumers who are most likely to respond to the brand's advertising are on the steepest part of the curve, in the middle. We set the thresholds for these consumers at 20% and 80% (i.e., consumers with a 20% to 80% probability of purchasing the brand) and call them the **movable middles.**

Figure 2a shows how much the size of the movable middles group can vary by brand, and figure 2b shows how the movable middles are on the steep part of the marketing response curve, and are therefore more responsive to marketing effects. Figure 2c illustrates the same phenomenon differently: by plotting a consumer's anticipated change (lift) in purchase probability (due to marketing exposure) against their initial purchase probability, we can see that the lift is more impactful for people in the middle of the curve.

LOYALTY DISTRIBUTION BY MARKET SHARE

FIGURE 2A: Size of the movable middles group for three hypothetical brands with different market shares.

MARKETING RESPONSE CURVE

FIGURE 2B: The movable middles are on the steeper part of the curve and are therefore more responsive to marketing effects.

MARKETING RESPONSE CURVE

FIGURE 2C: The movable middles experience the highest lift in purchase probability after marketing exposure, compared to people who buy the brand rarely or often.

7 Guadagni, Little "A Logit Model of Brand Choice Calibrated on Scanner Data", 1983 *Marketing Science* Vol. 2, Issue 3, pp 203–238
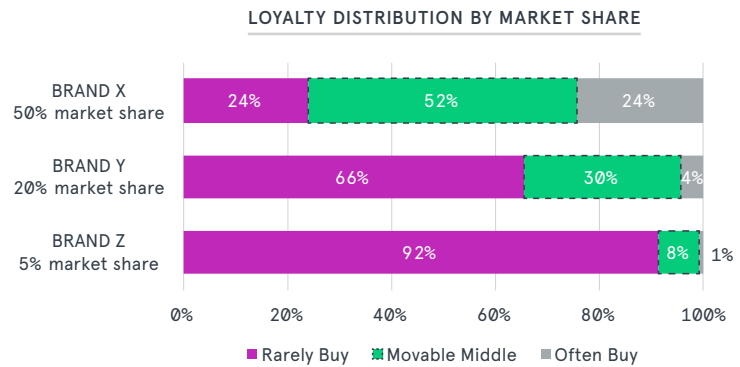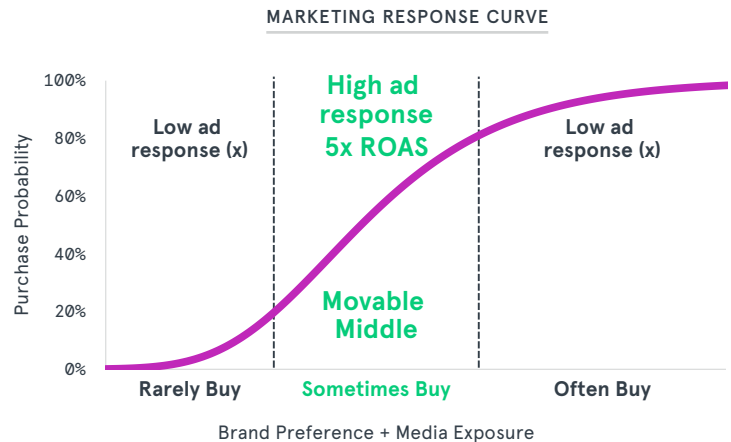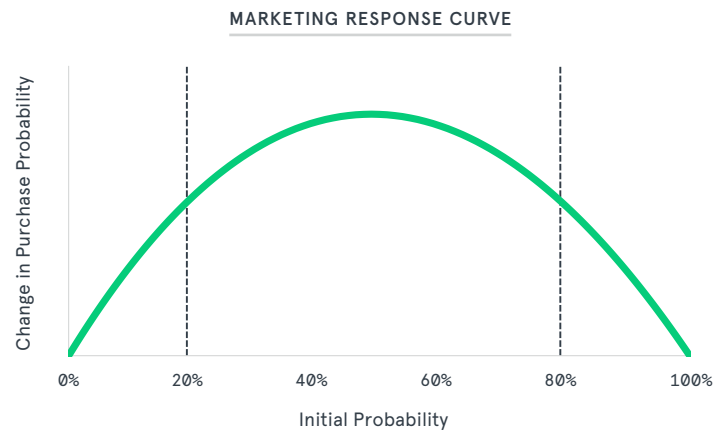
Key Processes

Agent-Based Simulation

We used an agent-based simulation to construct a sample of 623K 'synthetic' consumers for this project, representative of all US households active in the product category.

Why use agent-based modeling (ABM) and not real-life consumer data? We wanted to have the flexibility to simulate multiple markets with varying numbers of brands, market shares, repeat-buying patterns, ad responsiveness and media allocation scenarios without having to contend with the setbacks and compromises that come with real-life data (e.g., incomplete data, or bias due to prior media exposure). We used real-life data for calibration purposes and to validate the suitability of our models, but the comparative analyses were performed on ABM data.

ABM is a proven technique used not just in marketing and economics to perform simulations and predictive analytics, but in ecology, epidemiology, social networking and other fields to explore how individual behavior, governed by simple rules, can create complex ecosystems and reveal new patterns.

Real-Life Consumer Data

We validated key assumptions throughout the project with an academic and industry Advisory Board,[8] and used real-life data from Numerator[9] to calibrate our agent-based simulation and check the fit of our probability models. Figure 3 illustrates how closely our ABM approach matched Numerator data on a variety of measures.

While the bulk of the results we're presenting in this paper relate to a brand of frozen pizza,[10] we also used Numerator data to validate our model selection for 46 brands across 4 separate categories (frozen pizza, laundry detergents, nutrition bars, and butter/margarine). The correlation between simulated and real-life data was over 90% for all brands, and frequently over 99%.

MEASURE	SIMULATED DATA	NUMERATOR DATA
Market share (of trips)	10.3%	10.2%
Penetration (12-month)	25.0%	26.2%
Repeat rate	45.0%	44.0%

FIGURE 3: Close match between simulated and real-life data for brand A

MMA | neustar

Segmentation

Audience segmentation has a long history in marketing.

Human beings are complex thinking machines, to twist a term from economist Herb Simon:[11] each of us has a unique set of needs, beliefs, values and interests, and the brands we buy are an expression of our unique and complex personality. Yet it took a long time for marketers to realize that demographics and simple socioeconomic factors are insufficient to explain the diversity of consumer behavior.

Thankfully today, audience segmentation is ubiquitous. But no two segmentation schemes are the same. We selected the ElementOne (E1)[12] segmentation system from Neustar for this analysis. At the outset, we used our ABM simulation to generate 623K synthetic agents (i.e., consumers) and grouped them into 172 E1 representative segments based on shared demographic and life stage factors.

When we ran our beta distribution model for the study brand through each of those segments, some segments (like segment #166,[13] for instance) showed a high proportion of movable middles, while others (like segment #2[14]) had very few movable middles.

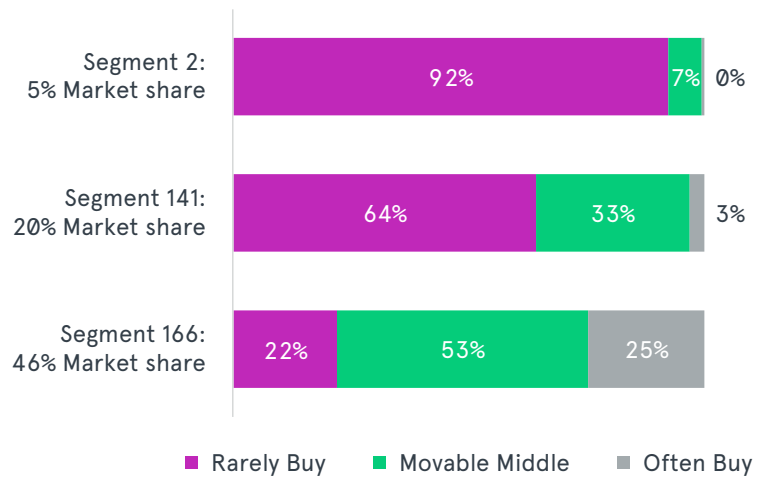**LOYALTY DISTRIBUTION BY MARKET SHARE**

FIGURE 4: Proportion of movable middles for select E1 segments for simulated brand

Demographics and simple socioeconomic factors are insufficient to explain the diversity of consumer behavior.

Media Allocation Plans

With a hypothetical $10 million campaign budget and eight weeks of airing, we developed three candidate media plans for this project:

Base media plan

The base plan was validated by our Advisory Group and developed to reflect current practices, CPMs and standard media allocation trends across channels. Figure 5 shows the details of this baseline plan.

MEDIA CHANNEL	BASE PLAN ALLOCATION	CPM	CAMPAIGN LENGTH	IMPRESSIONS (MILLIONS)	AVERAGE FREQUENCY
Linear TV	35%	$28.00	8 weeks	125	2.0
Desktop Video	15%	$20.30	8 weeks	74	2.5
Mobile Video	20%	$16.10	8 weeks	124	3.3
Desktop Display	10%	$6.15	8 weeks	163	4.7
Mobile Display	20%	$6.00	8 weeks	333	8.6

FIGURE 5: baseline media plan

Reach media plan

The reach media plan was created by calculating the reach elasticity of each media channel (adding and removing $500K at a time and measuring the marginal impact of that additional investment), and then using that elasticity to reallocate ad spend between channels based on what channel would increase unduplicated reach for the lowest effective CPM.

To stay close to real-life media planning practice, exposure frequency was capped at five impressions per channel, per week, for all digital channels, except for TV to mimic real-life limitations.

Outcome-based media plan

The outcome-based media plan was designed to target the movable middles. For this plan, we first computed ROAS for each of the 172 E1 segments from the base plan and used that ROAS to rank all segments from top to bottom. We made a few adjustments to the channel allocation plan based on the relative lift in performance for each media (shifting dollars from TV to digital video, for instance) and added some rules to mimic real-life constraints (e.g., at most 70% of consumers were deemed targetable on digital channels, and 25% on linear TV).

All targetable dollars were then reallocated to the segments with the best ROAS,[15] and the remainder of the budget was served across all audiences based on their likelihood to consume each media. Like in the reach plan, we applied a frequency cap of five impressions per channel, per week, for all digital channels.

These steps allowed us to zero in on the segments with the greatest proportion of movable middles.

The following charts show how the outcome-based plan differed from the other plans in the way that it allocated the $10 million budget between marketing channels (fig. 6a) and across the ROAS-ranked segments grouped in quartiles (fig. 6b).

MARKETING INVESTMENT ALLOCATION BY CHANNEL

Linear TV · Desktop Video · Mobile Video · Desktop Display · Mobile Display
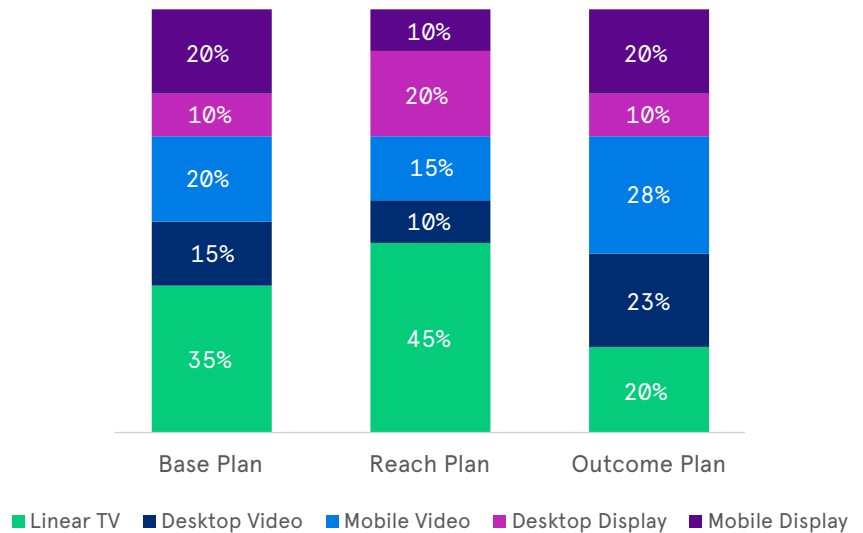
FIGURE 6A: Allocation of marketing funds by channel

MARKETING ALLOCATION BY ROAS QUARTILE

Top · 2nd · 3rd · Bottom

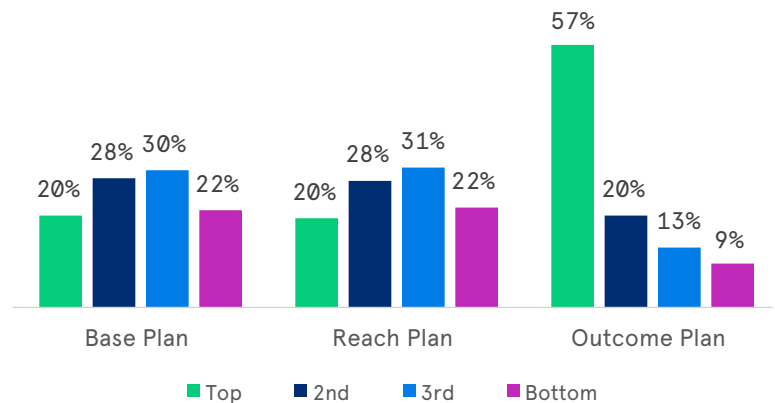**FIGURE 6B:** Allocation of marketing funds by segment quartile

15 A total of 49 segments had a ROAS over $3.

▪ Key Findings

1

The movable middles are a key new target group for a brand

By trimming off consumers who are very unlikely to buy the brand and those who are very likely to buy it (and therefore can do with less advertising), marketers can zero in on the consumers that will be most responsive to advertising. For a brand (brand A) of frozen pizza with a 10% market share, ROAS was close to 5X greater for its movable middles than it was for the rest of the population.
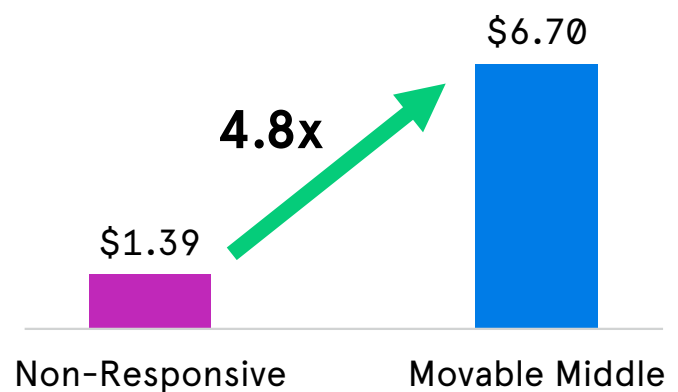

FIGURE 7: For brand A, ROAS is 4.8X greater for its movable middles vs. the rest of the population

MMA | neustar

2 The movable middles are a substantially larger group than heavy buyers

The movable middles are a substantially larger group than heavy buyers,[16] giving marketers interested in multiple outcomes a much larger target to work with.

Figure 8a shows that our brand of frozen pizza has more than twice as many movable middles than heavy buyers (16.7 vs. 7.7 million). Overall, the brand has 31.1 million yearly customers (one-fourth of all 120 million households in the U.S.), and the movable middles account for more than half of them (16.7 of 31.1 million). And figure 8b shows that 68% ($217 million / $318 million) of the brand's yearly sales can be attributed to them.

BRAND CUSTOMERS	PURCHASE PROBABILITY		
BUYER TYPE	**RARELY BUY**	**MOVABLE MIDDLE**	**OFTEN BUY**	**TOTAL CUSTOMERS**
Non-Buyers	2,720,196	1,017,598	61,800	**3,799,594**
Light Buyers	7,788,388	4,830,393	325,400	**12,944,180**
Medium Buyers	1,811,797	4,417,793	393,999	**6,623,590**
Heavy Buyers	418,399	6,461,190	852,199	**7,731,788**
Total Customers	**12,738,781**	**16,726,975**	**1,633,398**	**31,099,153**

FIGURE 8A: Number of brand A customers by purchase probability and buyer type

BRAND REVENUE	PURCHASE PROBABILITY		
BUYER TYPE	**RARELY BUY**	**MOVABLE MIDDLE**	**OFTEN BUY**	**TOTAL REVENUE**
Non-Buyers	$13,184,620	$5,809,527	$371,519	**$19,365,667**
Light Buyers	$38,564,582	$31,787,376	$2,337,980	**$72,689,938**
Medium Buyers	$17,099,398	$49,383,573	$4,607,705	**$71,090,676**
Heavy Buyers	$6,181,047	$129,994,651	$18,391,076	**$154,566,774**
Total Revenue	**$75,029,646**	**$216,975,127**	**$25,708,281**	**$317,713,054**

FIGURE 8B: Brand A revenues by purchase probability and buyer type

16 For the purpose of this analysis, light buyers are those who purchased the brand once over the eight months leading up to the start of the ad campaign, while medium buyers purchased it twice, and heavy buyers three or more times. As a result, the labels 'light,' 'medium' and 'heavy buyers' refer to purchase frequency.

The movable middles can be identified via modeling and targeted via segmentation

The crux of our method, for any given brand, is in identifying the segments with the largest concentration of movable middles. With that knowledge, marketers can drive higher performance through the careful reallocation of media dollars to those segments where their brand's movable middles are most prevalent.

An important step in our approach was to rank all segments in descending ROAS order. Figure 9 shows how that ranking came out for Brand A, and how the movable middles were distributed across those segments. The correlation is strong, and a media plan that zeroes in on the top-performing segments is likely to pay dividends.
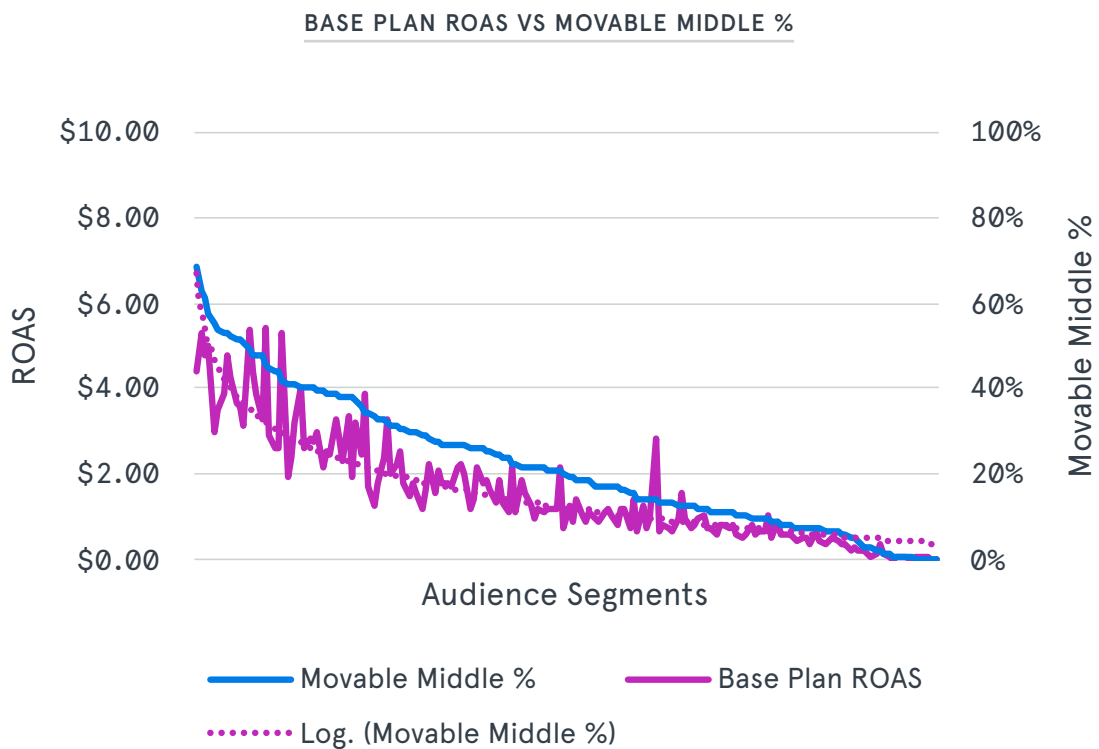
BASE PLAN ROAS VS MOVABLE MIDDLE %

FIGURE 9: Top E1 segments ranked by Movable Middle %

The outcome-based plan outperforms the reach plan by over 50% on ROAS while reaching more consumers and achieving higher penetration

The outcome-based media plan outperformed the reach plan by more than 50% on ROAS.

Figure 10a compares the performance of the two plans over the course of the media campaign, and figure 10b shows that the outcome-based plan, over the course of the media campaign, reached more consumers than the reach plan, across all buying groups (from non-buyers to heavy buyers).[17] The outcome-based plan is a clear winner on both counts.[18]

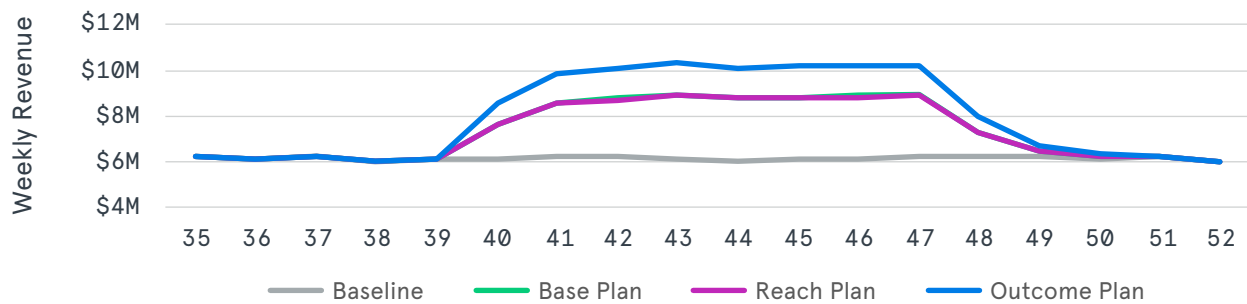| MARKETING PLAN	TOTAL REVENUE	CAMPAIGN REVENUE	REVENUE LIFT	ROAS	ROAS % INCREASE
Base Plan	$101.5M	$21.8M		$2.18	
Reach Plan	$101.2M	$21.5M	($0.3M)	$2.15	(1.3%)
Outcome Plan	**$112.5M**	**$32.7M**	**$11.2M**	**$3.27**	**50.2%**

FIGURE 10A: Comparative performance on ROAS for the duration of the campaign: the outcome-based plan outperforms the reach plan.

PENETRATION BY BUYER TYPE FROM BEGINNING OF CAMPAIGN TO END OF THE YEAR

BUYER TYPE	MEDIA PLAN		
	BASE PLAN	REACH PLAN	OUTCOME PLAN
Non-Buyers	5.6%	5.6%	**6.3%**
Light Buyers	31.1%	31.1%	**33.0%**
Medium Buyers	43.5%	43.6%	**45.7%**
Heavy Buyers	57.3%	57.2%	**59.3%**
Overall	**13.5%**	**13.5%**	**14.4%**

FIGURE 10B: Comparative performance on cumulative market penetration: the outcome-based plan outpaces the reach plan across the board.

17 For the purpose of this analysis, light buyers are those who purchased the brand once over the eight months leading up to the start of the ad campaign, while medium buyers purchased it twice, and heavy buyers three or more times. As a result, the labels 'light,' 'medium' and 'heavy buyers' refer to purchase frequency.

18 Targeting isn't free, but even if we consider that it might add $1-3 to the CPM in some channels, the boost in ROAS remains substantial.

5

The outcome-based plan reaches receptive new buyers

An outcome-based media plan can be used to grow the brand by attracting non-buyers (and light buyers) who are nonetheless open-minded about the brand. Figure 11a shows that ROAS among non-buyers in the top segments is much greater than it is for the bottom segments, and figure 11b shows how much better it is than for non-buyers reached by the reach plan.

By targeting segments where consumers are most likely to be in the movable middles group, the outcome-based plan is also finding 'lookalike' consumers who may not have bought the brand before but are more likely to respond to the brand's marketing than other consumers. Outcome-based marketing is a much more effective way to grow the brand than an indiscriminate reach plan would be.
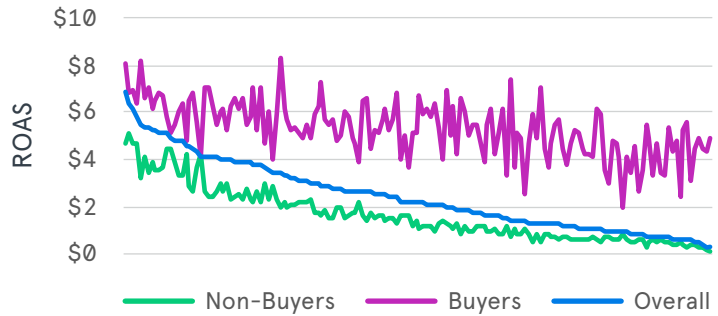
FIGURE 11A: Non-buyers in top performing segments show greater ROAS

NON-BUYER ROAS

$2.22

1.9x

$1.19

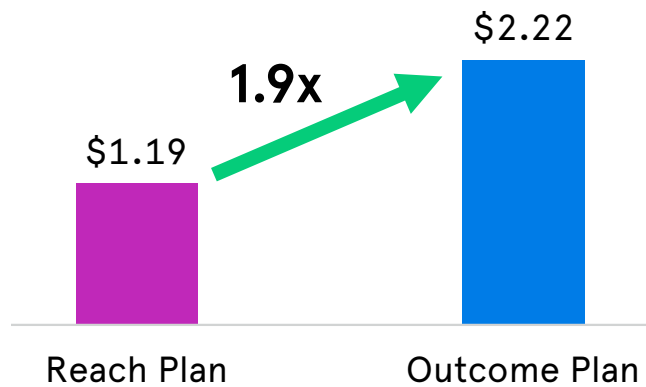Reach Plan   Outcome Plan

FIGURE 11B: ROAS for non-buyers in the outcome-based plan vs. reach plan

Outcome-based marketing is a much more effective way to grow the brand than an indiscriminate reach plan.

Next Steps

In an effort to simplify our analysis, we made some assumptions that should form the basis for further research.

For instance, we assumed that the advertising creative in all channels was of average effectiveness. Creatives are a crucial part of any marketing campaign, and it would be interesting to see the effect of creative quality on the various outcomes discussed in this paper. For instance, it's reasonable to imagine that a targeting campaign, like the outcome-based plan, would take this opportunity to customize creatives and tailor messaging to the movable middles. The improvements discussed in this paper should be amplified.

Another thread for future research pertains to advertising attentiveness. Different consumers have different levels of attentiveness for creatives on different channels. An outcome-based campaign that targets the movable middles with a marketing mix that maximizes that attentiveness would most likely outperform our current model. It also assumes that responsiveness across media is inversely related to the cost of that media. By

incorporating actual marketing lift by media channel and segment for a given brand, and optimizing across media channels for each segment, we would expect to see further opportunities for performance improvement.

Our approach was developed for mature brands in a stationary market, but we believe that this framework can be applied to emerging brands and markets. With some adjustments in the modeling functions, it would also be interesting to see how a similar approach might help marketers operating in markets where consumer interactions are more discrete (i.e., non-continuous) and contractual in nature (like magazine subscriptions or mobile plans).

Finally, we accounted for the short-term effects of our media plans, not the long-term effects, but we speculate that an outcome-based approach would be beneficial to the brand beyond the 12-week window of our research study. In fact, the outcome-based plan is built, at least partly, to attract more non-buyers to the brand, and those new buyers have more expected lifetime value for the brand ahead of them.

Acknowledgments

This work is the result of a collaboration between the Mobile Marketing Association and Neustar. We would also like to thank the following people for serving on our Advisory Board and steering this research project in the right direction:

Marketer Advisory Board

Bill Cramblit, Senior Manager Marketing Analytics & Optimization, Molson Coors

Tim Lewis, Director, Customer Acquisition, Choice Hotels International

Kevin Moeller, Head of Media Insights & Analytics, PepsiCo North America

Kay Vizon, Director of Media, Kroger

Academic Advisory Board

Dr. Dominique Hanssens, UCLA Graduate School of Management

Dr. Andrew Stephen, University of Oxford

Dr. Felipe Thomaz, University of Oxford

Subject-Matter Expertise

Joel Rubinson,
President Rubinson Partners, Inc and MTA subject matter expert for the MMA

Marketing Science Partners

Neil Morley, VP Data Science, Neustar

Marc Vermut, VP Marketing Solutions, Neustar

Appendix: Classic Models of Purchase Behavior

Customer-base probability models have been the backbone of academic research for decades. They've helped the industry understand important measures like loyalty, churn, and customer lifetime value, using relatively simple input variables linked to purchase frequency and recency.

To understand brand choice in a competitive market, two leading modeling approaches rose to prominence in the 1970s. One of those approaches was introduced by the Hendry Corporation,[19] and the other was spearheaded by Andrew Ehrenberg at the London Business School at the time. Although the mathematics of both systems were related at the start, they wound up going in two completely opposite directions in terms of implications.

The Hendry theory is based on modeling the distribution of purchase probabilities for consumers of a given brand, and it studies how that distribution leads to that brand having a certain market share in the product category—and how it reaches the appropriate level of repeat purchases to sustain that market share.

The primary probability model that emerged from this work was the Beta distribution,[20] which is the model we used in this paper. From a marketing recommendations point of view, the Hendry theory focused on brand loyalty, and the extent to which increasing loyalty was instrumental to increasing market share.

The model developed by Andrew Ehrenberg et al. took the focus away from brand choice for a single brand, and studied instead the distribution of purchases for all brands in a given category over time. The key breakthrough of that approach was the adoption of the NBD-Dirichlet distribution.[21]

The Dirichlet distribution is a multinomial beta distribution, modeling the distribution of probabilities of purchase of all brands simultaneously. The NBD (Negative Binomial Distribution) part of the model translates these probabilities into streams of purchases within a given timeframe. In the process, however, the Dirichlet model makes the unrealistic assumption that there is no market structure, whereas Hendry applications have found that market structure (usually based on product feature sets) does exist in virtually all product categories.

From a practical point of view, the NBD-Dirichlet approach led to a vision of the world where the key driver to increase a brand's market share is penetration: not repeat purchases among loyal customers, like in the Hendry approach, but increased penetration among category buyers. In the NBD-Dirichlet view espoused today by the Ehrenberg-Bass Institute, there are no weak or strong brands, only big or small brands. And broad reach, rather than targeting, is the preferred way to grow a brand.

So, is it more productive for a brand to maximize reach, or to target those who are most likely to respond to the brand's advertising? That's the central question we set out to address in this paper.

19 Kalwani et al. "A Parsimonious Description of the Hendry Model," Management Science, 1977, Vol. 23, Issue 5
20 Stewart "The Beta Distribution as a Model of Behavior in Consumer Goods Markets" Management Science, Vol. 25, No. 9, 1979, pp. 813–821
21 Goodhardt, et al. "The Dirichlet: A Comprehensive Model of Buying Behaviour," Journal of the Royal Statistical Society, 147 (part 5), 1984, pp. 621–55

ABOUT MMA

Comprised of over 800 member companies globally and 14 regional offices, the MMA is the only mobile trade association that brings together the full ecosystem of marketers, tech providers and sellers working collaboratively to shape the future. The MMA's global headquarters are located in New York with regional operations in Asia Pacific (APAC), Europe/Middle East/Africa (EMEA), and Latin America (LATAM).

ABOUT NEUSTAR

Neustar is an information services and technology company and a leader in identity resolution providing the data and technology that enables trusted connections between companies and people at the moments that matter most. Neustar offers industry-leading solutions in Marketing, Risk, Communications, and Security that responsibly connect data on people, devices and locations, continuously corroborated through billions of transactions. Neustar serves more than 8,000 clients worldwide, including 60 of the Fortune 100. Learn how your company can benefit from the power of trusted connections here: www.home.neustar.