

Morphing Theory and Applications

by

Gui B. Liberali

John R. Hauser

and

Glen L. Urban

May 2016

forthcoming, *Handbook of Marketing Decision Models (New Edition, 2017)*

International Series in Operations Research & Management Science published by
Springer Science and Business Media, Berend Wierenga and Ralf van der Lans, Editors

Guilherme (Gui) B. Liberali is Associate Professor of Marketing, Rotterdam School of Management, Erasmus University, Rotterdam, The Netherlands, (31)6.2096.4125, liberali@rsm.nl

John R. Hauser is the Kirin Professor of Marketing, MIT Sloan School of Management, Massachusetts Institute of Technology, E62-538, 77 Massachusetts Avenue, Cambridge, MA 02142, (617) 253-2929, hauser@mit.edu.

Glen L. Urban is the David Austin Professor in Management, Emeritus, Professor of Marketing, Emeritus, Dean Emeritus, Chairman of the MIT Center for Digital Business,, MIT Sloan School of Management, Massachusetts Institute of Technology, E62-536, 77 Massachusetts Avenue, Cambridge, MA 02142, (617) 253-6615, glurban@mit.edu.

The final publication will be available at Springer. <http://doi.org/978-3-319-56941-3>

1. The Morphing Concept

As electronic commerce often matches or exceeds traditional bricks-and-mortar commerce, firms seek to optimize their online marketing efforts. When feasible, these firms customize marketing efforts to the needs and desires of individual consumers, thereby increasing click-through-rates (CTR) and conversion (sales). When done well, such customization enhances consumer relationships and builds trust.

A/B testing is a popular means to optimize marketing efforts. The firm compares two or more communications vehicles, say two banner advertisement or two website implementations. For example, potential consumers (website visitors) are randomly assigned to two banners—one might emphasize general brand image and one might emphasize the comparative advantage of a product's features. The firm measures response in the form of CTRs or conversion to identify the better banner. The better banner is then used in day-to-day website operations. A/B testing can be used with multiple marketing instruments or with aspects of marketing instruments that are mixed and matched in an experimental design. A/B testing has proven effective and has increased the profitability of many marketing instruments.

Morphing improves A/B testing in many ways. First, morphing uses optimal adaptive experimentation. For example, as the morphing system begins to observe consumer response it allocates sample to A versus B to learn efficiently. Morphing trades off learning about consumer response (learn) with using that knowledge to display the best banner for the consumers (earn). The learn-while-earning process allocates same to different banners to maximize long-term profits. For example, if a morphing system learns that a particular banner is unlikely to be the best banner, it ceases to assign consumers to see that banner. If a morphing system learns that a particular banner is especially promising it automatically and optimally allocates more consumers to that banner.

Second, morphing automatically identifies the latent segment to which each consumer belongs. Morphing detects a consumer's segment from the clicks that the consumer makes on the firm's website (or from tracking the consumer prior to visiting the firm's website). For example, a consumer with a more-verbal cognitive style might

click more often on text-based descriptions than on pictures, whereas a consumer with a more-visual cognitive style might click more often on pictures. Alternatively, a consumer who is beginning his or her search for automobiles might click on comparison charts while a consumer who is ready to buy might click on dealer-location or special-deal pull-down menus.

Third, morphing matches marketing instruments to each consumer's segment, and does so optimally. Because morphing identifies latent segments automatically, morphing can use optimal experimentation for each segment to learn the best marketing instrument for that segment. For example, if the consumer has a verbal cognitive style, then the look and feel of the website can "morph" to feature more verbal content. If the consumer is in the buying stage for an automobile, then the website can help the consumer find dealers or cars with specific features. It might even offer an incentive for a test drive.

Fourth, because morphing identifies the best marketing instrument for each segment from those that are tried, it provides rich information for further development and design of those instruments. Indeed, in our experience, this organizational learning has proven to be critical to enhanced outcomes for the firm.

1.1 Morphing Overview

In this chapter we review almost ten years of morphing experience. To date, most of the contributions have been proof-of-concept research projects, but, increasingly, firms are beginning to adopt and test morphing capabilities. We begin with a brief overview on the steps in a prototypical morphing application

Morphing, as first proposed by Hauser, Urban, Liberali, and Braun (HULB, 2009), consists of the following steps:

1. Clicks on a website are monitored and, from those clicks, algorithms automatically infer the likelihood of the consumer segment to which the consumer belongs.
 - a. Websites, banners, or other marketing materials may be designed so that they appeal (potentially) to different segments of consumers.

- b. Consumers in a calibration study visit example websites and provide data by which to identify their segments.
 - c. The calibration data provides a model of how browsing behavior differs by segment.
2. Marketing materials, such as banners, are provided to consumers to maximize goals such as profit, sales, or click-through rates.
 - a. The system learns as is goes.
 - b. Learning is automatic and near optimal.
 - c. The goal of learning is to match the marketing materials to the consumer's segment to maximize the firm's goals.

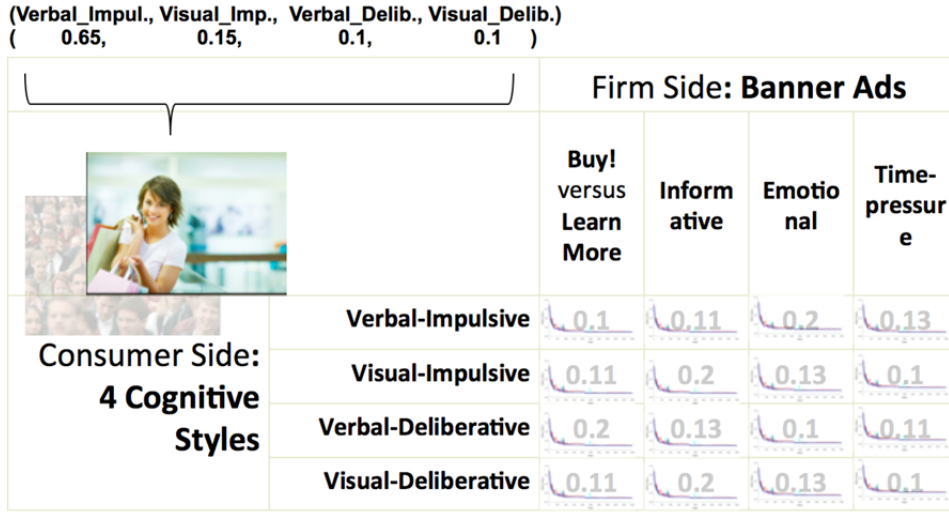
Ideally, morphing targets and learns from each and every consumer and does so in real time. However, some systems now “batch” learning in the sense that the rules in the second step are updated periodically rather than for every consumer (e.g., Bertsimas and Mersereau 2007; Schwartz, Bradlow, and Fader 2016). Recently, morphing has been extended to automatically determine when is the best time to morph the marketing materials (e.g., Hauser, Liberali, and Urban 2014).

1.2 Morphing Example

Figure 1 illustrates the general concept of morphing with a stylized example from banner advertising. This concept is not limited to banner advertising; morphing applies to a wide range of marketing materials. For example, HULB morph the look and feel of the website.

The set of numbers in the upper left of Figure 1 are the firm's best guess at the segment to which the consumer belongs. These estimates are based on the consumer's clickstream up to this point. For example, the firm might believe that this consumer's cognitive style is most likely to be verbal-impulsive (65%), but there is a lesser chance that the cognitive style might be visual-impulsive (15%), verbal-deliberative (10%), or visual-deliberative (10%). The morphing algorithm uses Bayes Theorem to estimate these probabilities from the clicks that the consumer has already made on the website.

*Figure 1: The Morphing Concept-An Example with Four Styles and Four Morphs
(Stylized Illustration)*



We provide details on the method later; we provide here the intuition. Suppose that, in an earlier calibration study, we measured consumer’s cognitive styles using traditional methods. For example, we might ask the consumer to answer a banks of questions, the answers to which indicate the consumer’s cognitive style. Suppose further that we observed that consumers with verbal-impulsive cognitive styles clicked on action-oriented textboxes, but consumers with other cognitive styles clicked on other portions of the website. Then if we observe that consumer clicks on many textboxes and prefers short action-oriented descriptions rather than longer fact-based descriptions, that consumer may more likely to be verbal-impulsive than one of the other cognitive styles. The actual probability then is proportional to the percent of consumers with verbal-impulsive cognitive styles (known from the calibration study) times the likelihood that a person who clicks on action-oriented verbal textboxes is verbal-impulsive (also known from the priming study). Because a website is likely to offer a large number of click choices, we describe each click by its characteristics to reduce dimensionality. The end result, which is updated when the consumer provides more clicks, are the percentages in the upper left corner of Figure 1.

Before we describe the learn-while-earning aspect of morphing, it is easier if we consider a more traditional situation in which we have observed a large number of consumers of each cognitive style for each of the potential banners. For this situation, the

table in the lower right side of Figure 1 contains outcome probabilities for each banner-segment combination. For example, if a consumer has a cognitive style that is “verbal-impulsive” and that consumer is given the “Buy! verses Learn More” banner, the likelihood that that consumer would click on the banner is 0.10. If that consumer were given instead the “Emotional” banner, then the likelihood of a click would increase to 0.20. These probabilities are based on prior consumers, with those cognitive styles, who have been shown each of the banners. If this were an A/B test, the banners would have been randomly assigned until we had sufficient precision on the outcome-probability estimates.

If the firm had perfect information about the segment to which the consumer belonged and if it knew the outcome probability perfectly, the firm would select the banner with the highest outcome probability by looking up the highest outcome probability in the row corresponding to the consumer’s segment. For example, if the firm knew the consumer was visual-impulsive and it knew the outcome probabilities, it would provide the consumer with the “Emotional” banner because it has the highest outcome probability in the verbal-impulsive row.

However, firms do not know the consumer’s segment with certainty. Instead, based on the consumer’s clicks, firms has estimates of the likelihood that the consumer belongs to each of the four cognitive-style segments. If the system had completed its learning, the best banner would be the banner that maximizes the firm’s immediate goals such as CTR. In Figure 1, the best banner for the consumer is the Emotional (in column 3) because it has the highest expected reward given our estimate of the consumer’s cognitive style. We obtain that estimate by multiplying the probabilities the consumer belongs to each cognitive-style segment times the probabilities that a consumer in that segment clicks through when shown a given banner. For example, the likelihood that the consumer clicks through when given the Emotional banner is 0.17 obtained as $0.17 = 0.65 \times 0.20 + 0.15 \times 0.13 + 0.10 \times 0.10 + 0.10 \times 0.13$.

Emotional is the best banner to provide to the consumer after the system has completed its learning, but Emotional may or may not be the best banner to provide to the consumer while the system is still learning the outcome probabilities. For example, it

might be best if the system were to occasionally try other banners so that it can learn probabilities for other banners. Furthermore, the system might be able to make the best decision even if it does not know the outcome probabilities with certainty. For example, it might be able to eliminate banners with extremely low outcome probabilities and not waste sample consumers on those banners. This is a dilemma. If the firm tries other banners it sacrifices its goals for the immediate consumer, but the system learns what is best for the next consumer and all subsequent consumers with the same cognitive style. The morphing system automatically assigns banners to consumers with (near) optimal experimentation by balancing the cost of experimentation with the value of learning about the outcome probabilities.

The mathematics used to balance earning and learning are sophisticated, but the implementation is relatively straightforward using a mathematical concept called “Gittins’ indices (GIs).” At any given time there is a GI for each cognitive-style-banner combination. Each GI represents both the value the firm can gain from that consumer (the current best estimate of the outcome probability) and the option value to the firm for learning more about outcome probabilities. We provide more details later. Gittins (1979) proved that there is a rule based on GIs that provide optimal experimentation when consumers can be assigned to consumer segments without error. The rule is simple, provide the morph with the highest GI. Because the rule is optimal, a firm using GIs can expect higher profits than it would earn with naïve A/B testing.

However, as Figure 1 indicates, we do not know the consumer’s segment with certainty. But we have probabilities, based on the consumer’s clicks, that the consumer belongs to a segment—e.g., a 65% chance that the consumer’s segment is verbal-impulsive. Krishnamurthy and Mickova (1999) demonstrated that if one were to use GIs rather than the best estimates of outcome probabilities, then the multiple-latent-segment experimentation would be near optimal. This is exactly what is done in morphing. We replace the outcome probabilities with GIs and compute the expected value—the expected Gittins’ index (EGI). We assign to the consumer the banner with the highest EGI.

Morphing has an additional advantage over standard A/B testing. Because the

system is always in learn-while-earning mode, the GIs automatically and near optimally pick up any changes in the underlying outcome probabilities. For example, if consumers' tastes change and a banner is no longer as effective, the GIs will begin to shift automatically to take into account the option value of learning more about that banner. Similarly, if a new banner is added to the mix, the GIs for that banner start at prior beliefs, but the morphing system quickly and optimally learns the true outcome probability.

Morphing uses information from prior consumers to learn the updated value of the outcome probabilities for each banner-segment combination. This value is updated after each user is exposed to a banner—after we observe whether or not the consumer clicks-through or makes a purchase (conversion). The value of each banner-segment combination GI is based on our current estimate of the choice probability for that banner-segment combination plus an option value that reflects the value of learning more about that banner-segment combination. As the system learns from many observations, the option value decreases. For the banner that is best for a consumer segment, the value of the banner-segment combination converges to the predicted purchase value. For other banner-segment combinations, the system might cease to allocate sample because it is not profitable to do so, even considering the option value of learning.

The rate at which convergence is achieved is an empirical question, and is based on the expected traffic of the website. For example, it is optimal for websites with millions of visitors per day to learn at a different rate than websites with only a few thousand visitors per day. In practice, the degree to which each individual observation changes beliefs about the best morph depends on how many observations we expect to observe during the relevant time period.

Morphing is not limited to banners – it can be used to match consumers to website designs, call-center scripts, or any marketing instrument. For the remainder of this chapter, when it is clear in context, we use the general term “morph” instead of banner to reflect the generality of potential applications.

Morphing is based on continuously improving initial knowledge about each consumer's segment, using that information to assign the best morph for a consumer, and

learning from the outcome of the morph assignment. As shown in the next section, morphing has been applied in a variety of situations, including applications to improve performance for firms which had previously run randomized controlled experiments (A/B tests).

2. Optimal Online Experimentation

2.1 From Learn-Then-Earn to Learn-While-Earning: The Multi-Armed Bandit Problem

Randomized controlled experiments are the cornerstone of causal inference. Firms run hundreds, or even thousands, of online randomized experiments every day, in what is often referred to as A/B testing. Typically, an A/B test is based on random assignment of treatments to website visitors and is continued until sufficient statistical power is achieved such that reliable conclusions can be made regarding the effect of a specific website configuration on sales or other variables of interest. For example, one firm may run an online A/B test to learn whether it is more effective to present information about its product in a 2- or 3-column format. The dependent variable is typically CTR or online purchase (conversion).

Because of randomization and statistical power, traditional A/B tests tend to follow the learn-then-earn paradigm. During the testing phase the focus is on *learn*, i.e., estimating the effect of each treatment on consumer behavior. Once the estimates are obtained, the focus changes to outcome maximization (*earn*), when the firm deploys the winning treatment on large scale.

The traditional *learn-then-earn* paradigm of A/B testing has two major weaknesses. First, it is based on the responses of average consumers; it ignores heterogeneity in consumer preferences. It does not take into account that different consumers may respond differently to the marketing instrument(s) when the firm deploys the winning treatment (marketing instrument) to all consumers. When consumers are not all the same, ignoring individual differences can be costly. For example, in Figure 1 it is better that the verbal-impulsive consumer get an Emotional banner while the visual-deliberative consumer get an Informative banner.

Morphing addresses this issue using consumer segment probabilities (for latent

consumer segments) to handle heterogeneity when computing the optimal treatment for each consumer. Morphing enables each consumer (or each segment) to get the best morph based on the latest information about the behavior of the consumer's segment. "Best" takes into account both learning and earning.

Second, the learning phase in traditional A/B tests is inefficient, leading to wasted resources because it invests the same amount of resources on good and bad treatments. Typically, A/B tests assign the same sample size to each cell during the learning phase, which means that the precision of the estimates of good treatments is the same as the precision of the estimates of the bad treatments. However, as the firm learns quickly that a treatment is suboptimal, it wastes resources when it assigns more consumers to a suboptimal morph in order to make its estimate of the outcome probability more precise for that treatment.

Morphing invests sample size in those cells that most clarify which marketing instruments to give to which consumers. Because traditional A/B testing continues to invest sample in learning about suboptimal treatments, the firm loses revenue every time a treatment is assigned to a cell that the firm already knows has a low probability of leading to a good outcome (click or conversion). In morphing, once the firm is confident that a marketing instrument is best for a consumer segment, it optimally assigns that marketing instrument to almost all subsequent consumers in the segment.

Solving this learn-while-earning problem is not easy. Obtaining better estimates about the effect of marketing instruments (*learn*) is costly in the short-term, but leads to higher revenue on the long term. On the other hand, using current estimates to assign marketing instruments to consumers (*earn*) avoids the short-term cost of learning, but suffers from higher opportunity costs. The firm misses future sales because it does not learn which marketing instrument is really best for each consumer segment. For example, if there is no exploration, then, if current estimates suggest that a 3-column design does not have the highest conversion rate for a specific segment, the 3-column design will never be shown again to any consumers in that segment. This loss of future potential can loom large, particularly if some "shock" changes outcome probabilities.

The learn-while-earning problem is at the heart of morphing. This problem is in

the class of “multi-armed bandit” problems. When segments are known, website morphing methods provide an optimal solution to this problem in real-time (HULB). When segments must be inferred, the solution is not provably optimal, but is extremely close to optimal. Morphing dynamically –and near optimally - allocates larger sample to the best treatment-segment combination based on the solution of the learn-vs.-earning formulation originally developed by Gittins (1979).

2.2 From Learning about Designs to Learning about Consumers

Adapting a website to each consumer involves a fundamental change in the philosophy of A/B testing. Typically, A/B testing assigns banners or website variations to consumers on a random basis. As a result, an A/B test identifies the marketing instrument that is best on average, not the marketing instrument that is tailored to each consumer. In some cases, the marketing instrument might be best for no one. For example, suppose that consumers are either Type X or Type Y and suppose there are three morphs, A, B, and C. Suppose that the outcome probabilities for Type X consumers are 0.9, 0.5, and 0.0 for A, B, and C, respectively. Suppose they are 0.0, 0.5, and 0.9 for Type Y consumers. On average, the best morph is B, with average outcome probability 0.5. However, if we could assign A to Type X consumers and B to Type Y consumers, we would achieve an improved outcome probability of 0.9. Customization matters.

Morphing changes the A/B testing logic fundamentally. Instead of testing marketing instruments that apply to all consumers, morphing learns and selects the best marketing instrument for each consumer. Instead of randomly assigning marketing instruments to a test or control treatment, morphing optimally assigns consumers to marketing instruments. As more and more consumers are run through a morphing system, the algorithm identifies the best allocation of morphs to consumers to maximize the outcome variable (such as conversion).

Changing A/B test focus from A vs. B marketing instruments to a focus on consumers is a major shift for most firms doing A/B testing. The change in focus has two practical implications. First, most firms and A/B software assign incoming consumers randomly to test or control cells. The software is not designed to learn about consumers

and then assign consumers to different marketing instruments based on information about that consumer and the accumulated experience from other consumers.

Second, morphing requires tracking consumer-level information. Most large firms today use software packages that act as a layer isolating managers from the raw data. Reports are produced automatically with summary statistics showing which marketing instrument is the best on average, and at which p -value. Obtaining reports based on individual consumers (instead of marketing instruments) requires access to and analysis of raw data, something that is often a formidable task for most firms.

Morphing requires firms to change radically the way they design and run their A/B tests, and the way they use information about website visitors (consumers). In our illustration in Figure 1, we defined consumer segments by cognitive styles. This is illustrative only. We can define consumer segments by the stage in the buying process, interest in the category, cultural styles, cognitive styles, source (whether the consumer is coming from an online search or a referral), personas, devices (tablet, desktop/laptop or mobile), purchase tendencies, or any other variable that can be observed in a calibration study.

2.3 Handling Consumer Differences: The Case for Cognitive Styles

Although consumer segments can be defined in a variety of ways, one of the most frequent ways to segment website consumers is based on the way they interact with websites and other morphs. The way consumers respond to websites is heavily related to how they gather, process, and evaluate information—their *cognitive styles* (Hayes and Allinson 1998). A cognitive style reflects “individual differences in how we perceive, think, solve problems, learn and relate to others (Witkin, Moore, Goodenough and Cox 1977, p. 15).” Examples of dimensions of cognitive styles include impulsive-deliberative, visual-verbal, and analytical-holistic (for more examples, please refer to the online appendix in HULB).

If measured well, a consumer’s cognitive style is stable over time, so there are no history-dependent interactions (Markovian structure) which would make it difficult for the morphing algorithm to converge to true outcome probabilities. Decades of research in psychology suggest that people develop cognitive styles over the years, and that their

preferences for cognitive styles change slowly.

Cognitive styles are easily interpretable and actionable from a managerial point of view. Designers can easily relate to cognitive styles, such as verbal-visual, to develop website designs, banners, or other marketing instruments that are likely to be suited well for one style rather than others. The more a morph is tailored to a specific style, the more likely a consumer using that cognitive style will relate to and feel comfortable with the way the website, banner, or other marketing instrument communicates with the consumer. Strong prior beliefs help morphing converge faster, but, even if the designers' prior beliefs are wrong, morphing learns optimally the best morph-to-segment assignments based on consumer reaction. The GIs converge automatically to the best assignments based on consumer response even if the initial designers guess incorrectly.

3. Learning Loops: Consumer Segments and Morph x Segment Assignments

For ease of exposition we illustrate morphing with an application based on cognitive styles.

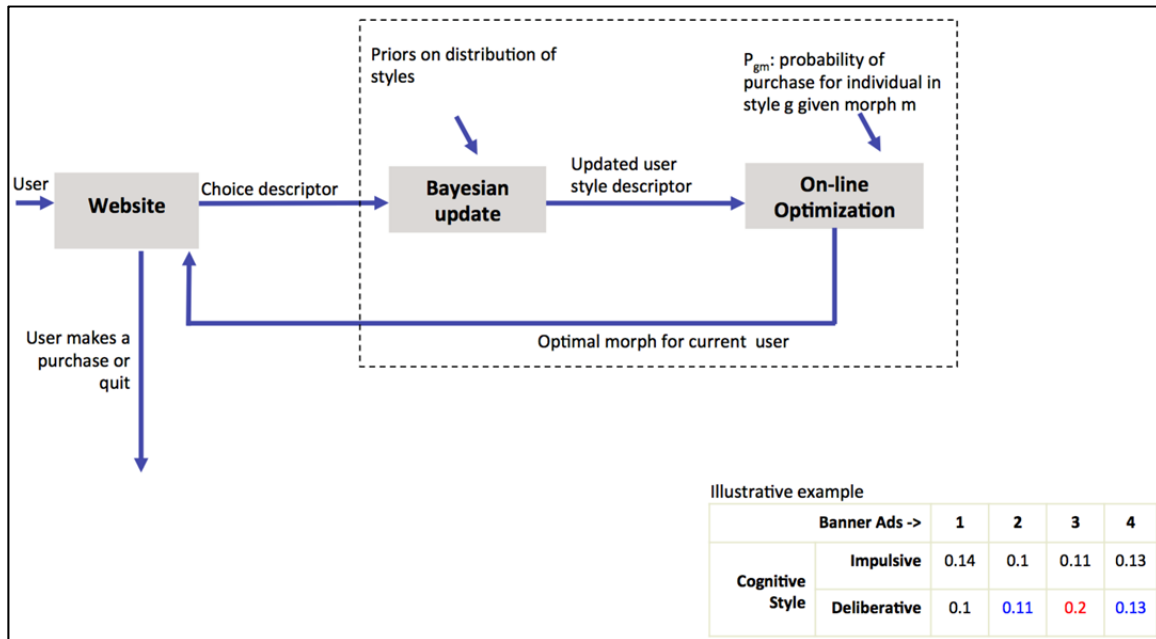
The morphing process has two learning loops. In the first learning loop, the morphing system observes clicks from each consumer and, after sufficiently many clicks, updates its estimates of the probability that consumer belongs to each segment. In earlier versions (Morphing 1.0), the number of clicks was set exogenously. Later versions (Morphing 2.0) choose the number of clicks endogenously and near optimally for every consumer.

In the second learning loop, the system learns the outcome probabilities *across consumers* in a segment as those consumers are exposed to morphs (treatments) and respond with successes (such as a click through or a purchase conversion) or failures (a non-click-through or a non-purchase). Figure 2 illustrates the process. For ease of exposition we reduced the number of cognitive styles from four in Figure 1 to two in Figure 2. The basic concepts apply to as many cognitive styles as can be defined, but when more cognitive styles are used, more data (consumers) are needed for the morphing system to work well.

The process starts when a consumer comes to the website. At this point (typically,

on a landing page), the consumer may be exposed to a morph-independent stimulus. We observe the first click (or set of clicks). A click (or set of clicks) can be thought of as a choice among various links, all of which have cognitive cues. When the consumer chooses one of the links, we gain data about the consumer’s cognitive style. For example, one consumer may decide to go to the virtual advisor area of a website by clicking on a verbal description instead of clicking on an image. Using this information about the consumer’s click choices, the learning-about-the-consumer loop updates prior beliefs about the consumers’ cognitive style using the Bayes Theorem. The resulting posterior beliefs become the updated probabilities that the consumer has either an impulsive or a deliberative cognitive style.

Figure 2: Morphing and Learning



Next, the morphing system uses the information about the consumer’s cognitive style to look-up the GIs in an optimality table, as in the lower right of Figure 2. Recall that the GIs indicate the value of each morph (highest earn-vs.-learn value) for each segment. The GIs are larger than the expected outcome probabilities because they include the option value of learning. As we did for Figure 1, we use the probabilities that the consumer belongs to each segment to compute an expected Gittins index. We choose the morph with the highest expected Gittins index, breaking ties randomly. The consumer is

exposed to the morph with the highest Gittins index. For example, suppose that the learning-about-the-consumer loop predicts that there is an 80% probability that the consumer is deliberative. Then the expected Gittins' index, EGI, for the third banner advertisement would be $0.00352 = (0.20)(0.11) + (0.80)(0.20)$. This is higher than the EGI for Banner 1 (0.00224), Banner 2 (0.00176), or Banner 4 (0.002704).

The consumer continues to browse until he or she either clicks-through (if CTR is the outcome measure) or buys (if conversion is the outcome measure) or leaves the website without clicking-through or buying. Based on the observation of the consumer's response to the third banner advertisement, the optimality table is updated accordingly as described in the next section. (Because there was still some uncertainty in identifying the consumer's segment, the response probability for both the impulsive and deliberative cognitive styles would be updated, albeit the deliberative style more so than the impulsive style. There would be no updates for the first, second, and fourth banners.)

Recall that the GIs are not outcome probabilities. If we were to assign morphs to consumers based on outcome probabilities, we would assign the morph with the largest expected outcome. Such a rule would never assign any other morphs and we would never improve our knowledge about outcome probabilities associated with those morphs. This non-assignment problem is known as the *curse of serendipity*, or lack of exploration.

The informal intuition behind Gittins' solution to this dilemma is that the GIs summarize the value of earning-vs.-learning as an optimality index. In particular, a GI equals our estimate of the outcome probability augmented by the value of exploration. A GI is computed for every cognitive-style-morph combination and allows the system to explore and serve morphs to learn how to assign the best morph for each consumer. As morphs are served and outcomes observed, we 'spend' the exploration value. When all exploration value is spent, only the true outcome probability remains in the updated table. §4 presents a more formal description of this approach. An appendix provides the analytical expressions used to implement the expected Gittins index solution.

3.1 Steps in a Morphing Project

Prior to implementing the morphing algorithm in day-to-day operations, parameters must be estimated in a calibration study. See Figure 3. It is also feasible to use

4. The Analytics of Morphing

The Morphing 1.0 algorithm was first published in HULB. The Bayesian updating effects a rapid assessment of consumers to segments. A dynamic program produces the GIs to select the best morph for a segment, An expectation over the GIs identifies the best morph. The Morphing 2.0 algorithm was published in Hauser, Liberali, and Urban (2014). In this section we provide the basic morphing concepts. An appendix provides the formal notation and equations.

4.1 Learning the Consumer's Segment from the Consumer's Clicks

The Bayesian model in morphing was motivated by a Bayesian advisor that identified which vehicle (car or truck) to recommend to a consumer (Urban and Hauser 2004). The recommendation was based on that consumer's answers to a series of questions about the potential consumer's use of the vehicle. In our case, we use the consumer's clicks on the website to identify the likelihood that the consumer belongs to a particular consumer segment rather than to identify the best recommendation.

Let n index consumers, r index segments, m index morphs, and t index clicks for each consumer. At any point in the consumer's visit to the website, the consumer has a choice of which click to make next. We characterize the probability of any click for a consumer with a particular cognitive style. To do this, we describe each click by a set of characteristics. For example, we might observe basic dimensions such as graphical vs. verbal, functional characteristics such as "use an analytic tool" or "read a post," or website areas such as "virtual advisor" or "learning center." HULB use eleven website characteristics for a broadband-sales website.

In the calibration study, morphing analysis estimates a choice model that includes weights for each characteristic. Because we know the characteristics of all possible clicks every time a consumer makes a click, we compute the "utility" of a given click as a function of the to-be-estimated weights for the characteristics. A logit model assumes that the consumer maximizes his or her "utility" as given by the click characteristics and an error term. We use the calibration data and standard maximum-likelihood or Bayesian methods to estimate the weights.

After we estimate the weights, we use these weights in day-to-day website operation to compute the probability, for each cognitive style, that the consumer will choose a given click. In stylized symbols, we compute: $Prob\{click \mid cognitive\ style\}$. The likelihood of a particular clickstream is just the product of these probabilities multiplied for all the clicks made by the consumer. Equation 1 in the appendix provides details of the logit model likelihood. Note that, although we observe the clicks, we still need to compute the probability, as conditioned on cognitive style, for the next step in morphing.

From the calibration study, or from the history of consumers visiting the website, we form prior beliefs about the segment to which the consumer belongs. For example, we might believe that 25% of all consumers are verbal-impulsive. Call this probability, $Prob\{cognitive\ style\}$. We want to compute the probability of a cognitive style based on the observed click stream. We do this with Bayes Theorem recognizing that

$$Prob\{cognitive\ style \mid clickstream\} \propto Prob\{clickstream \mid cognitive\ style\} * Prob\{cognitive\ style\}$$

Where \propto means proportional. To compute the actual probability we normalize the expression so that $Prob\{cognitive\ style \mid clickstream\}$ adds up to 1.0 when summed over cognitive styles. Equation 2 in the appendix provides details. Fortunately, Equation 2 involves relatively fast calculations so that cognitive styles can be determined almost instantaneously between clicks on a website

4.2 Learning How to Assign Morphs to Segments Optimally

We represent our knowledge about outcome probabilities by a function that we can update quickly. In particular, we choose a “beta distribution.” The beta distribution has two parameters that depend upon the consumer’s segment, r , and the morph, m . These parameters are α_{rm} and β_{rm} . If p_{rm} is the probability that a consumer in segment r , who was shown morph m , clicks through (or converts), then, for the beta distribution, the mean outcome probability is $E[p_{rm}] = \alpha_{rm}/(\alpha_{rm} + \beta_{rm})$. Larger values of the parameters mean less uncertainty in our beliefs about the outcome probabilities.

The beta distribution allows fast updating. For example, if the cognitive style

where known, then α_{rm} increases by 1.0 for every success (click-through or conversion) and β_{rm} increases by 1.0 every time the consumer leaves without a “success,” e.g., without a click through or a conversion.

Assigning the optimal morph to the n^{th} consumer is more complicated than simply maximizing the immediate reward. The GI includes the option value. It does more than simply maximize $E[p_{rm}]$. Because each outcome improves our knowledge about p_{rm} , the updated distribution enables us to make better decisions in the future. The dynamic decision problem balances immediate rewards with the knowledge gained that enables better decisions in the future.

This dynamic problem for this type of multi-armed bandit was first formulated in the 1940s and, for many years, considered to have no simple solution. However, in the late 1970s, John Gittins had a seminal insight that he could compare the decision problem for each “arm of a bandit problem” to an equivalent fixed outcome. He could then compare the equivalent fixed outcomes from many arms and choose the outcome that was best. The value of the fixed outcome became known as a Gittins’ index. The concept was generalized to many problems. Today, if a problem can be solved with indices, it is said to be indexable. When cognitive styles are known, the morphing problem is indexable.

The basic dynamic program is formulated as a recursion known as a Bellman equation. The “state” of the problem is the current values of α_{rm} and β_{rm} , as well as a discount factor, a . The discount factor indicates how much the morphing algorithm should discount the future. For example, if a website has 100,000 visitors spread equally throughout the year, then HULB suggest that $a = 0.999999$.

The recursion recognizes that, for any given rm combination, the best strategy is to choose the larger of the fixed outcome or to keep experimenting. Gittins’ proved that once the fixed outcome is chosen, the best strategy is to continue choosing the fixed outcome. (This makes sense, α_{rm} and β_{rm} are not changing if there is no experimentation. When there is no experimentation on an arm, choice among alternatives does not change and neither does the solution.) If the algorithm chooses to try that morph for that cognitive style, then we get to observe an outcome—either a success or a failure. But we know the likelihood of a success, $\bar{p}_{rm} = E[p_{rm}]$.

With probability, \bar{p}_{rm} , we observe a success and reap our reward, which we set to 1.0. With probability $1 - \bar{p}_{rm}$, we observe a failure and reap no reward. In each case we get to continue playing the game with the updated α_{rm} and β_{rm} . Equation 3 in the appendix provides the details. We provide here the recursion in words. Let V indicate the value of continuing to play with a given set of parameters. Then, the recursion is:

$$V(\text{current } \alpha_{rmn}, \beta_{rmn}, a) = \max \left\{ \frac{GI_{rmn}}{1-a}, \bar{p}_{rmn}(1 + V(\text{updated based on success})) + (1 - \bar{p}_{rmn})V(\text{updated based on failure}) \right\}$$

The first term reflects the value of continuing to $t = \infty$ with a discount factor of a . In this equation we added the subscript n to indicate that these values change after each consumer. To use the recursion, we solve the equation with an iterative search for all of the α 's and β 's we expect in practice. We table the GI's so that the GI's can be assessed quickly.

When a consumer's segment is not known with certainty, the dynamic program becomes a partially-observable Markov decision process (POMDP). In general POMDPs are difficult to solve, but this particularly POMDP has a near-optimal solution that runs in real time between clicks (Krishnamurthy and Mickova 1999). Specifically, the revised algorithm replaces the Gittins' index with the expected Gittins' index, EGI, as illustrated in §3. Equation 4 in the appendix provides the details.

Finally, we update beliefs about the parameters of the beta distribution. The challenge is that we do not know the consumer's latent segment with certainty. We only know the probabilities that the consumer belongs to each of the consumer segments. The true Bayesian updating formulae are no longer easy, but we can use a trick. When the segments are latent, we can update if we consider "fractional observations" using an analogy to the standard likelihood function. (Fractional updates represent a pseudo-Bayesian updating that provides estimates that work extremely well for morphing. See formal analyses and simulations in Hauser, Liberali, and Urban 2014.)

If we observe a success, conditioned on the consumer having seen morph m , we consider this as a fractional success for each latent consumer segment, r . The fractional

of the success is probability that the consumer was in consumer segment r . The binomial distribution is well-defined for fractional observations and naturally conjugate to the beta distribution, so we use the same formulae, except with fractional observations. Equation 5 in the appendix provides the details. Updating occurs when the consumer leaves the website. The fractional-observation formulae enable the morphing algorithm to run in real time between a consumer's clicks on the website.

4.3 When do We Know Enough to Find the Optimal Morph?

In Morphing 1.0, the algorithm morphed after a fixed number of clicks by the consumer on the website. For example, in HULB's application to a BT Group website that sold broadband service, the a morph was considered after the 10th click. In Urban, et al. (2014)'s application to banner advertising on CNET, the banners were morphed after the 5th click. In both cases, the time to morph was set by the researchers' judgment based on simulated performance. We can do better by choosing the click on which to morph.

In choosing the time to morph, we address the tradeoff between exposure and precision. We gain greater exposure of the best morph to the consumer by presenting the optimal morph as early as possible in the consumer's website visit. Doing so exposes the consumer to the best morph for the longest amount of time possible. We gain greater precision by identifying the best morph as late as possible in the consumer's website visit, because doing so uses better consumer-segment estimates to find the best morph. To address this trade-off, Morphing 2.0 uses a second recursion.

The generalized morphing algorithm, published in Hauser et al., (2014), solves an embedded dynamic program that enables the optimal trade-off between exposure and precision for each consumer. To formulate the dynamic program, the authors had to first address three issues in consumer response to morphing.

First, to evaluate the impact of every morph that the consumer sees, the algorithm must explicitly track how long a consumer is exposed to each morph, and decide how to attribute credit to each morph seen. Second, the algorithm must allow the system to change morphs as often as necessary because, as more information on the cognitive style becomes available from clicks, beliefs about the true cognitive style become closer to the

true cognitive style. This introduces memory into the algorithm because the algorithm must keep track of how many clicks were made for each morph exposure for the consumers. This challenge recognizes that the optimal morph after many clicks may be different from what was thought to be the optimal morph when the algorithm had less information about the consumer cognitive style (fewer clicks). Third, if the algorithm allows multiple morph changes and allows the system to decide when to morph, consumers might experience cognitive load from seeing multiple morph. This cognitive load induces potential cognitive switching costs that must be modeled. By modeling switching costs, the algorithm only changes morphs if the gains from changing morphs are greater than the cognitive costs of switching morphs.

4.3.1 Every Morph Seen Matters: the Attribution Problem

If we allow a consumer to be exposed to more than one morph during a website visit(s), we need to attribute credit regarding the observation (a success or a failure) to each morph seen. For example, assume a consumer saw Morph A during the first five clicks, then saw Morph B during the last ten clicks, and then made a purchase. Which morph should get the lion's share of the credit for this success? Should it be the first morph because "first impression lasts?" Or, should it be the second morph because it was seen for longer, or perhaps because of recency effects?

We address this attribution problem by specifying attribution weights, w_t 's, for each time period, t , when computing value of a morph for consumer n . The weights, w_t 's, are measured, judged, or estimated empirically in each application, and used as parameters of the model. Because w_t is applied to each morph seen at every time t , it spreads the credit through all morphs seen. To keep the number of w_t 's small, we allow t to index observation periods that may be one click or more than one click. We normalize the impact weights so that they sum to 1.0 over clicks (or observation periods).

4.3.2 Changing Morphs: Switching Costs

It is reasonable to expect that consumers may experience cognitive load if the website design changes too often or too dramatically. The costs of switching tasks have been extensively studied in psychology starting with Jersild (1927) and Spector and

Biederman (1976), and more recently in Meiran (2000). In marketing, switching costs are well-established (e.g., Weiss and Anderson, 1992; Jones, Mothersbaugh, and Beatty 2000, 2002). Researchers have also studied how consumers react to switching costs when browsing websites (Balabanis, Reynolds and Simintiras, 2006, and Johnson, Bellman and Lohse , 2003).

Additive switching cost are common in the multi-armed bandit literature and algorithms exist (e.g., Banks and Sundaram 1994; Dushochet and Hongler 2003; Jun 2004), but additive switching costs require that we keep track of the timing of all switches for a consumer. This path dependence makes it more difficult to solve the optimization problem. Because of this difficulty, additive switching costs make algorithms infeasible for real-time morphing.

On the other hand, a multiplicative switching cost can be factored out in a recursive equation that optimizes the time to morph. Multiplicative switching costs are more intuitive because their effect is proportional to the likelihood of purchase. Not only do multiplicative switch costs assure that all probabilities remain defined between zero and one, but we expect that the amount by which a low probability is lowered by a switching cost would be less than the amount by which a high probability is lowered by a switching cost. For example, suppose a switch lowers p_{rnm} from 0.800 to 0.700. Comparable proportional cost would lower p_{rnm} from 0.090 to 0.070 while a comparable additive cost would lower p_{rnm} from 0.090 to less than 0.000. To date, we have not tested the multiplicative assumption, but it seems to be a more-reasonable representation of switching costs than an additive assumption.

For both practical and theoretical reasons, we solve a problem with multiplicative switching costs and do so in real time. Specifically, we assume that a switch in a morphs lowers the consumer's purchase probability. The switch lowers the purchase probability by a factor of γ where $\gamma \leq 1$. In theory, γ can be determined by experimental means in a priming study. However, to date, γ has been set by managerial judgment. Hauser, Liberali, and Urban (2014) explore the sensitivity of γ between 0.80 and 1.00.


4.3.3 Putting it Together

The tuning parameters, the switching factor (γ) and the period-weights (w_t 's), must be selected before the algorithm is used to morph a website (in day-to-day operations). The tuning parameters, require either managerial judgment or experiments during the calibration study. In a calibration study, segment membership is measured directly, therefore the true consumer segment is known among calibration respondents. To estimate tuning parameters, the calibration study would also assign switches randomly at different time periods. With sufficiently many observations in the calibration study, γ and the w_t 's can be identified.

Figure 4 illustrates the conceptual decision problem for the case where the consumer makes a purchase (or leaves the website) after four observation periods. Specifically, during the first observation period, the website displays a morph. The respondent makes clicks while exploring the website and we update our beliefs about the consumer's segment, Using the new information, and anticipating more information from subsequent decision periods, we decide which morph to display in the second decision period if the consumer stays on the website. (The consumer may decide to leave after a decision period. For example, the consumer might leave after the t^{th} observation period with probability, ψ_t .) If the morph in the second period is different from the morph in the first period, the consumer incurs a multiplicative switch cost, γ .

This process continues until the consumer reaches the fourth period at which time the consumer either purchases or leaves without purchasing. Figure 4 illustrates the process as if the consumer makes a decision after the fourth period. However, in practice, the consumer can make a decision at any period and/or continue beyond the fourth period. (The recursive equation in the appendix allows random exits.)

Figure 4: The When-to-Morph Problem (modified from Hauser, et al. 2014)



In general, the when-to-morph decision problem is coupled with the learn-while-earning decision problem where the learn-while-earning algorithm experiments with different morphs for each segment. For example, if we show morph m to a consumer in the consumer segment r for more data periods, we learn more about the response probability for that segment-morph combination (p_{rm}). Fortunately, the dynamics of the two decision problems happen on two very different scales. The “which-morph-for-which-segment” learn-while-earning decision problem is solved from observations based on success over failure over thousands of consumers. On the other hand, the when-to-morph decision problem is solved between clicks for one consumer at a time.

Because of these differing dynamics we decouple the two problems. In particular, we use the Gittins’ indices (GIs) to represent the value of showing morph m to a consumer in segment r and we use the concept of the expected Gittins’ index (EGI) to decide the best morph. The GIs are updated between consumers and are held constant when the when-to-morph decision problem that is solved between clicks by the current consumer.

Putting this altogether we obtain a recursive relationship that must be solved

between clicks for the current consumers. Unfortunately, this recursive relationship does not appear to be indexable. The conceptual recursive equation is the following, where t indicates the observation period. To keep this recursion simple, we have not specified random exit. Equation 6 in the appendix provides greater details.

$$V_t(m_t^*, m_{t-1}, clicks) = \max_{m_t} \left\{ w_t EGI + \sum_s Prob\{segment = s\} V_{t+1}(m_{t+1}^*, m_t, clicks | segment = s) \right\}$$

The equations that we used in Morphing 1.0 assumed that only the last morph seen by the consumer affects the probability of a successful outcome for that consumer. When we solve the when-to-morph dynamic program, we must generalize the fractional observation updating procedure. In particular, we keep track of which morph was shown to the consumer in each observation period. The fractional observation is now the probability (based on all observed clicks for that consumer) modified by the w_t 's. For example, if the consumer saw morph m_1 for the first period and morph m_2 for the second, third, and fourth periods, then, if the w_t 's are normalized to 1.0, the fraction assigned to morph m_1 for segment r is the (terminal) probability that the consumer is in segment r times w_1 . The fraction assigned to morph m_2 for segment r is the (terminal) probability that the consumer is in segment r times $w_2 + w_3 + w_4$. Equation 7 in the appendix provides details.

4.4 On-going Extensions and Other Methods

Shortly after the HULB was published, morphing was extended to handle longitudinal interventions. The extended algorithm was tested in a field experiment matching AT&T banner advertising with cognitive styles identified from clicks consumers made on CNET.com. See Urban et al., (2014). Other methods have been published addressing the application of multi-armed bandit ideas to morphing. Table 1 summarizes a few of these applications.

In general, we see that some methods, such as Thompson Sampling, focus on aggregate or batched data and on non-consumer-specific marketing instruments. Interestingly, Schwartz, et al. (2016) test alternative multi-armed bandit solutions in

counterfactual synthetic-data experiments and suggest that Gittins-based strategies often outperform Thompson-sampling-based strategies even in batched applications. Other heuristics also do well. The website morphing papers and Chung et al. (2009) learn at the level of the individual consumer to enable the system to match marketing instruments to consumers efficiently.

Table 1: Examples of Multi-Armed Bandit Algorithms for Online Experimentation

	<i>Focus (what is it learning about)</i>	<i>General or industry- specific</i>	<i>Considerations on Optimality</i>
<i>Hauser, et al. 2009, Hauser, Liberali, and Urban 2014</i>	<i>Consumer</i>	<i>General</i>	<i>Optimal for indexable problems (know consumer segments). Near-optimal for partially observable consumer segments.</i>
<i>Urban, et al. 2014</i>			
<i>Scott 2010, Schwartz, Bradlow, and Fader 2016</i>	<i>Creative</i>	<i>General</i>	<i>Designed to run in batches. Asymptotically optimal. Arms pulled proportionally to posterior probability of being optimal.</i>
<i>Bertsimas and Mersereau 2007</i>	<i>Creative</i>	<i>General</i>	<i>Designed to run in batches. Lagrangian decomposition and asymptotic approximations.</i>
<i>Chung et al., 2009, 2015</i>	<i>Consumer</i>	<i>Industry-specific</i>	<i>Promotes explorative search with a rejuvenation heuristic step</i>

5. Applications of Morphing: Evidence from the Field

The first application of morphing, as reported in HULB, was a research collaboration with the BT Group, formerly British Telecomm (BT). In this project, the data indicated that, had morphing been implemented system-wide, the lift in BT’s online sales of broadband plans would have increased by 20% – about \$80 million in additional revenue. These data were analyzed further in Hauser, Liberali, and Urban (2014). Their counterfactual synthetic-data experiments suggested that the improved Morphing 2.0 methods would have outperformed the original Morphing 1.0 methods by 69%. The gains

reflect the ability of Morphing 2.0 to handle switching costs, attribution, and the optimal time to morph. The Morphing 2.0 algorithm was applied to a website selling card loans in Japan. An initial study of 1,395 consumers provided data for counterfactual experiments that predicted a 63% improvement—substantially more than the Morphing 1.0 algorithm (Hauser, Liberali, and Urban 2014).

Urban, et al. (2014) applied morphing to banner advertising. In a field experiment on CNET.com, banner morphing achieved a 97% lift in click-through rates for context-matched banners relative to a no-morph control group. CNET.com is a high-traffic website that hosts display banner advertising; it was not feasible in the field experiment to track online sales (conversion). To examine the impact on conversion as well as click-through rates, Urban, et al. ran a longitudinal field study with General Motors' Chevrolet Division. The field experiment documented that matching banners to the stage of the consumer's buying process, body-type preference, and cognitive style significantly increased click-through rates, brand consideration, and purchase likelihood relative to a control.

We are aware of several morphing applications that are now being developed. For example, one application has begun a proof-of-concept test using traffic from a major telecomm provider in The Netherlands. The calibration study has been completed and the cognitive-style Bayesian loop has also been coded. This application includes four cognitive styles, three morphs, and several funnel-stage outcomes. It includes controls so that morphing can be evaluated.

A second morphing application is being developed in collaboration with an online marketplace in The Netherlands. This application morphs the automotive section of the online market place, and uses two consumer-knowledge segments instead of cognitive styles. The online marketplace is expanding its assignment mechanism to allocate consumers to test and control using the morphing algorithm. A third application is starting at a disruptive financial-products-comparison portal with operations in various countries in Asia. While none of these applications have yet gone live, they indicate the feasibility of developing morphing websites and morphing banners across a wide variety of applications.

6. Design and Implementation Decisions in a Morphing Project

Morphing methods substantially increase click-through and conversion rates because they fundamentally change the way website design, banner advertisements, and other marketing instruments are tested. Conversion managers and the IT teams involved in a morphing project benefit from the managerial and technical implications of such changes. This section provides an overview of key changes based on our experience with morphing projects in a variety of firms.

6.1 From Managing Aggregate Data to Handling Consumer-Level Data

Perhaps the most unexpected practical challenge companies face when considering morphing is the unprecedented need to handle data that is tagged to individual consumers. This operates on multiple levels.

- Morphing requires that the firm track and update its estimate of the probability that each consumer belongs to a segment. These updates might be based on data from the firm's website, but advanced applications base these updates on consumers' activities on many channels such as clicks on the website, posts on social media, or call-center input.
- Morphing requires that firms track, at least temporarily, which consumers are exposed to which morphs. Fortunately, the system needs only to maintain parameter updates and indices, not the entire morph-to-consumer history, but many websites must be modified to maintain even this level of information.
- Morphing requires that consumers be assigned to A/B cells dynamically. It is no longer sufficient to assign consumers randomly to A/B cells. Rather morphing bases these (near-optimal) assignments based on balancing immediate profit and long-term learning.
- Some firms may wish to test morphing itself versus a control such as random assignment or a fixed-morph control. In this case, random assignment of consumers to treatments occurs at a higher conceptual level. Consumers are assigned to strategies (morphing vs. a control) rather than marketing instruments. To the best of our knowledge, no off-the-shelf software has the capability of assigning consumers to strategies.

These four challenges require a major technological shift, because standard randomized-assignment code needs to be updated, reporting systems need to adapt, and firms need to rethink their policies on A/B testing. Morphing experiments identify which designs are best for which consumers, but, often, morphing also provides the organization with a new way to think about website, banner, and marketing-instrument design. Morphing provides a new way to manage click-through rates, conversion, and other funnel measures.

6.2 Consumer Segments, Marketing Instruments, and Outcome Variables.

Website morphing integrates three foundational elements of e-commerce: demand (e.g., consumer segments), supply (e.g., marketing instruments used by a firm), and online transactions (e.g., conversion, a request for information, or a click-through). This section provides an overview of what needs to be done for each element.

Consumer segments. While cognitive styles remain one of the best segmentation variables, morphing can be applied to a variety of segmentation variables including country of residence, personas, source (referral or not), device being used (mobile, computer, etc.), stage in the buying process, etc. The only real requirements are that consumers do not switch segments during a session and that segments can be identified in the calibration study.

Marketing instruments. Morphing can apply to any marketing instrument that can be tested with traditional A/B testing. Marketing instruments include website designs, banner advertising, call-center scripts, product recommendations, price levels, promotional coupons, etc. Marketing instruments (morphs) can also be defined at a higher level of abstractions, such as an advertising campaign that is implemented in several different online channels. For example, one morph for a telecomm firm could be a campaign focused on emotional content—the campaign might present its services as a way to keep close to family and friends. A second morph could be a campaign focused on informational content—the campaign might show how the quality of service is better than the competition. Both campaigns could run in parallel and be implemented in various media channels. Morphing would identify which segments of consumers relate best to which of the two campaigns. (Of course, in this case, consumers would need to be

tracked across channels.) Based on consumers' clickstreams, elements of the best campaign can be targeted to the right consumers. The elements might even be channel dependent.

Click-through rates, conversion rates, and other funnel measures. Click-through rates versus conversion rates are not conflicting goals, but they are often distant in time in the purchase funnel (e.g., Hongshuang and Kannan 2014). Analyses must be done carefully to infer causality when the temporal distance between exposure, consideration, and purchase is too long. Purchases in many categories do not happen immediately; purchases may happen weeks, or even months, after exposure to a marketing instrument. In the interim, there are often other changes in the website and in the environment that also affect sales. Either these changes must be modeled or managers must recognize the inherent uncertainty in end-of-the-funnel measures. Project leaders need to carefully clarify what are realistic optimization goals and then choose the outcome variable (funnel measure) accordingly. The choice of the outcome variable is crucial for the mechanics of morphing (what to optimize), for what firms learn, and for how the success of morphing is evaluated.

6.3 A Roadmap to Implement Morphing

Each morphing project has several tasks and milestones that need to be achieved.

1. Select the segmentation criteria, e.g., cognitive styles or other variables.
2. Select the morphs, e.g., marketing instruments such as website design or banners.
3. Select the outcome variable, e.g., click-through rates, awareness, trial, or conversion.
4. Determine the webpages and links to monitor. Perhaps design the webpages so that segments are easy to identify (next-generation websites).
5. Assess a categorization of each monitored link using a panel of judges for use in the Bayesian model (the website characteristics).
6. Run a calibration study to observe consumer's clicks and assign representative consumers to segments. In the calibration study consumers are assigned through direct questions.

7. Using the data from the calibration study, estimate a model that predicts click preferences for each consumer segment.
8. Pre-compute the click likelihoods, the probability of a click given the characteristics of the click (and competing clicks) and the consumer's segment. Do this for each link and each consumer segment so that segment likelihoods can be obtained quickly with the Bayesian model.

Coding

9. Implement the consumer inference system. This is the real-time Bayesian-loop inference code running on the webserver.
10. Define the control cell, e.g., random assignment, status-quo method, fixed-morph, or best-guess?
11. Decide whether to have a single test cell or multiple test cells. In a single-test-cell design, morphing chooses the optimal morph to show. One could potentially decide to run two or more test cells in parallel, each running a different set of morphs, segments, and/or outcome variables.
12. Adapt the existing A/B system to randomly assign (and keep track of) consumers to test and control cells.
13. Implement the system that receives the morph assignments and selects the morph to serve to the current consumer.
14. Adapt the reporting system to report consumer outcomes (click-throughs, conversions, or other funnel measures) based on the selected outcome variable. The system should report at the consumer-level for each morph the consumer received. A Morphing 2.0 system may also need to record the number of observation periods (sets of clicks) for each morph seen by the consumer.
15. Implement code that delivers the best morph to a consumer based on the morphing optimization.

6.4 Priors and Convergence

There are two sets of priors used to initialize the morphing system. The first prior represents initial beliefs about the consumer segment, before any click is observed. This

is typically selected to be either flat (equal probabilities to each consumer segment) or equal to the observed percentages of consumer segments in the calibration study. The decision depends on sample size, precision, and reliability of the estimates in the calibration study. It is relatively easy to update this prior after sufficiently many consumers have been observed in day-to-day operations. This prior is important because it affects which consumers get which morphs. Often, inferences about a consumer's segment must be made after a relatively few clicks by the consumer.

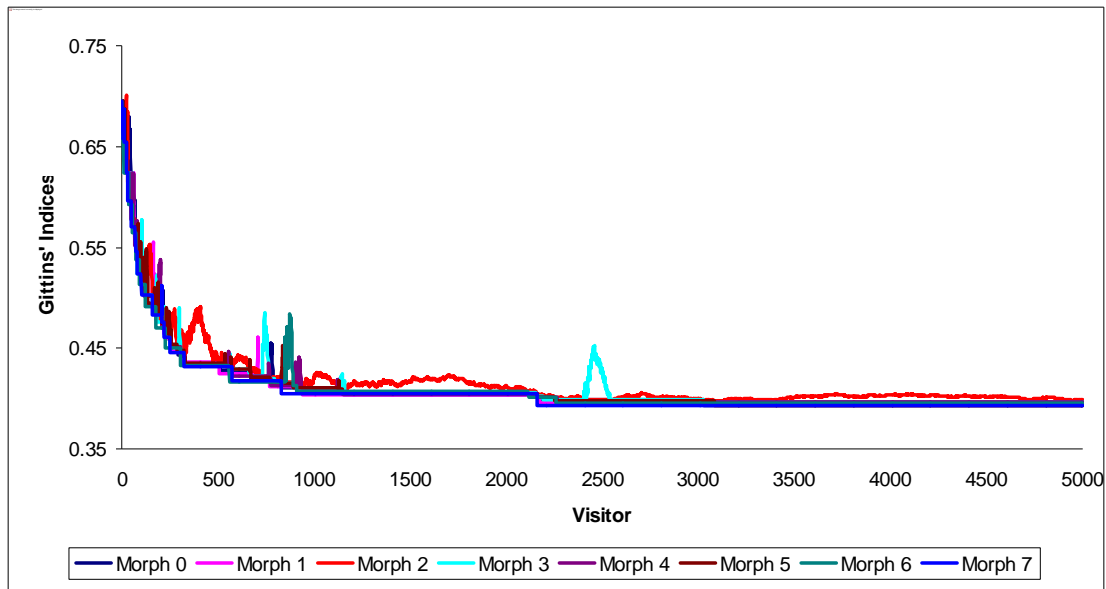
The second prior is the prior beliefs about the outcome probabilities. The Gittins' indices are calculated for the first customer based on priors that reflect the strength of our beliefs about the initial outcome probabilities for every morph-segment cell. Typically, the prior beliefs are based on observed outcomes of morph x segment probabilities in the calibration study. In some cases, the morphs may still be under development during the calibration study. In these cases, it is reasonable to start with flat priors (expected baseline click-through rate or another appropriate measure). Fortunately, in day-to-day operations, the performance of the morphing algorithm is relatively robust with respect to this prior on outcome probabilities. As consumers visit the website, the data on their clicks soon overwhelms the prior beliefs about outcome probabilities.

Typically we expect to see a pattern of transition from learning to earning that is somewhat similar to Figure 5. Figure 5 applies to a single consumer segment in a website that receives approximately 100,000 annual visitors. It documents how the GI for each morph changes as more consumers are exposed to morphs and as more outcomes are observed. Notice that, after a few thousand visitors, all indices have converged to the true morph x segment outcome probabilities. When the segment is not known but inferred using probabilities of the consumer belonging to a segment, convergence is not as rapid, but still occurs.

The rate at which the GIs converge is controlled by the discount factor, the volume of data, and accuracy of the cognitive-style posteriors. Applications with less concentrated cognitive-style posterior probabilities (probabilities that are closer to equally-likely) tend to converge more slowly than applications with concentrated probabilities (probabilities that are close to 0.0 or 1.0). Convergence is slower when

consumer segments must be inferred because fractional updating spreads outcome observations over multiple consumer segments. The right discount factor enables the GIs to be matched appropriately to the firm's cost of capital.

Figure 5: Convergence of Morphing Algorithm for a Cognitive Style (illustrative data from HULB)



7. Do's and Don'ts of Morphing and Organizational Impact

Although morphing makes morph assignments in real time while consumers click on websites, firms may decide to update the outcome probabilities for morph x segment combinations in batches. While this is possible, project leaders should be aware that it is only by multiple iterations of updates on that table (see Figure 1) that the system will learn optimally. In extreme cases, the learning loop can be run offline, but an offline learning loop is not an efficient use of resources. An offline system learns from past data, but does not realize gains from optimal experimentation. Similarly, if the Bayesian loop is run offline, the morphing system does not have the ability to match morphs to consumer segments.

Morphing tends to cross organizational silos in large, traditional corporations. A project typically requires efforts from the engineering team (to code the systems listed in the roadmap), sales teams (morphs are marketing instruments), web designers, reporting

(to develop the APIs that need to be integrated with cloud servers), and consumer experience teams. As in any project involving change, support from the highest level of the organization is crucial for success. If the firm does not have an established culture of A/B experimentation, morphing is likely to require additional steps. In such cases, our experience suggests that a pilot project can be implemented to optimize marketing-instrument A/B testing without the Bayesian loop, or to allocate marketing instruments to consumers without the second learning loop (the GIs). After A/B or consumer-level testing has been completed successfully, the full system will be easier to sell within the organization.

From a computational point of view, there are two major considerations. First, there is a need for real-time inference of consumer segments. Our formulae allow for rapid computation, but algorithms must be coded and implemented and may require re-training so that the web developers gain experience with Bayes-based algorithms. Second, the performance of the data transfer between the morphing servers (based on the cloud) and the firm's traditional webservers must be tested extensively to make sure performance is appropriate. Computations are designed to be rapid and the traffic that flows between servers is designed to be light (just a few bytes per click), but the connection between servers must have high levels of data reliability and speed.

From a purely methodological point of view, the use of Gittins' indices is based on a technical assumption – that the multi-armed bandit problem is indexable. Indexability in the canonical problem is usually expressed as a requirement that the arms of the multi-armed bandit are independent. For further details refer to Gittins et al. (2011). When the arms are not independent, the problem may still be indexable, but Gittins' indices may no longer be optimal. (When the “arms” change by external means, the multi-arm bandits are called restless bandits. Indices such as Whittle's index may need to be used—Whittle 1988.)

Morphing is flexible, but designers should be aware of a key tradeoff. The dimensionality of the optimality table grows proportionally to the number of morphs times the number of consumer segments. Successful applications balance relevance and speed. Simulations are valuable because they provide benchmarks before the firm

chooses the number of consumer segments and the number of morphs.

From the viewpoint of website development, webpages can be designed so that consumer segments are easy to identify. Links and other content can be planned in a way to maximize the information obtained from each click, reducing the number of clicks needed to learn the consumer's segment. We call such websites generation 2 (Gen-2) or next generation (next-Gen) websites.

8. Open Questions and Relevant Challenges

Morphing theory and methods provide opportunities for further research in substantive, conceptual, and methodological areas. From a substantive point of view, there are opportunities for new applications of morphing using other marketing instruments, such as price levels, promotion types, retention policies, call centers, and product bundles. Morphing is also feasible when using different devices or combinations of devices, namely desktop computers, tablets, and smartphones.

We are not aware of projects based on morphs built across media channels, but we believe that a consumer's clicks across online channels may substantially improve consumer-segment inference. There is effort to measure media across multiple channels (Liberali et al. 2015), but that does not include morphing. There are opportunities for morphing to coordinate actions that blend direct human action, such as calls, with automated actions, such as product recommendations. Because morphing affects organizational culture, there are opportunities to explore the feedback from morphing results to the creative process at agencies. Creative teams responsible for the development of website designs and banners obtain new creative insights by understanding which consumer segments respond best to which marketing instruments.

Morphing 2.0 provides a structure to model switching costs, attribution, and random exit. This structure opens opportunities in the measurement of switching costs, in the study of attribution, and in the modeling of exit probabilities. There are challenging issues in how to aggregate lessons learned over multiple A/B tests. Schwartz, et al. (2016) provide a means to address these issues within batch-processed A/B testing, but challenges remain for allocating marketing instruments to individual consumers in non-

batch modes. Advances in multi-armed bandit research provides many opportunities such as correlated-arms bandits (as in Keller and Oldale, 2003) and restless bandits (using the Whittle index as in Song, Zhang and Hauser, 2014).

See Urban et al. (2009) for more managerial issues. The online appendix of HULB provides additional insights on the development of cognitive styles and the appendices in Hauser, et al. (2014) provide details on a number of technical issues, including fractional updating.

References

- Bertsimas, D., Mersereau, A.J. (2007). A learning approach for interactive marketing to a customer segment. *Operations Research*, 55:6 1120-1135
- Chung, T.S., Rust, R.T., Wedel, M. (2009). My mobile music: An adaptive personalization system for digital audio players. *Marketing Science*, 28(1) 52-68.
- Chung, T.S., Wedel, M., Rust, R.T. (2015). An adaptive personalization using social networks. *Journal of the Academy of Marketing Science*, 42(1) 1-22
- Gittins, J.C. (1979). Bandit processes and dynamic allocation indices. *J. Roy. Statist. Soc. Ser. B*, 41(2) 148–177, plus commentary.
- Gittins, J.C., Glazebrook, K., Weber, R. (2011). *Multi-armed bandit allocation indices*. (John Wiley & Sons, London).
- Hayes, J., Allinson, C.W. (1998). Cognitive style and the theory and practice of individual and collective learning in organizations. *Human Relations*, 31:7 847-871.
- Hauser, J.R., Liberali, G., Urban, G.L. (2014). Website morphing 2.0: Switching costs, partial exposure, random exit, and when to morph. *Management Science*, 60(6) 1594 – 1616.
- Hauser, J.R., Urban, G.L., Liberali, G., Braun, M. (2009). Website morphing. *Marketing Science*, 28(2), 202-224.
- Hongshuang, L., Kannan, P.K. (2014). Attributing conversions in a multichannel online marketing environment: an empirical model and a field experiment, *Journal of Marketing Research*, 51, 40–56.
- Keller, G., Oldale, A. (2003). Branching bandits: a sequential search process with correlated pay-offs. *Journal of Economic Theory*, 113, 302–315
- Krishnamurthy, V., Mickova, J. (1999). Finite dimensional algorithms for the hidden Markov model multi-armed bandit problem. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 5 2865-2868.
- Liberali, G., Urban, G.L., Dellaert, B., Tucker, C.E., Bart, Y., Stremersch, S. (2015). A new method to measure online media advertising effectiveness: Prospective meta-

- analysis in Marketing. Working paper, MIT Sloan School of Management.
- Lin, S., Zhang, J., Hauser, J.R. (2014). Learning from experience, simply. *Marketing Science*, 34:1 1-19.
- Schwartz, E.M., Bradlow, E.T., Fader, P.S. (2016). Consumer acquisition via display advertising using multi-armed bandit experiments. forthcoming *Marketing Science*.
- Scott, S.L. (2010). A modern Bayesian look at the multi-armed bandit. *Applied Stochastic Models Business. Indust.*, 26(6):639–658.
- Urban, G.L., Hauser, J.R. (2004). “Listening-in” to find and explore new combinations of customer needs. *Journal of Marketing*, 68: 72-87.
- Urban, G.L., Hauser, J.R., Liberali, G., Braun, M., Sultan, F (2009). Morphing the web - building empathy, trust, and sales. *MIT Sloan Management Review*, 50(4): 53-61.
- Urban, G.L., Liberali, G., MacDonald, E., Bordley, R., Hauser, J.R. (2014). Morphing banner advertising. *Marketing Science*, 33(1): 27-46.
- Witkin, H.A., Moore, C., Goodenough, D., Cox, P. (1977). Field-dependent and field-independent cognitive styles and their educational implications. *Review of Educational Research*, 47 1-64.
- Whittle P (1988). Restless bandits: Activity allocation in a changing world. *Journal of Applied. Probability*, 25(2):287–298.

Acknowledgements

The on-going experiments are supported by the Erasmus Research Institute in Management (ERIM).

Appendix: The Equations of Morphing

Let n index consumers, r index segments, m index morphs, and t index clicks for each consumer. Capital letters, R , M , and T_n denote totals. Let c_{tn} denote the t^{th} click by n^{th} consumer and $\vec{c}_{tn} = \{c_{1n}, c_{2n}, \dots, c_{tn}\}$ denote the vector of clicks up to an including the t^{th} click. At each click choice, the consumer faces J_{tn} click alternatives as denoted by c_{tnj} where j indexes click alternatives. Let $c_{tnj} = 1$ if consumer n clicks the j^{th} click alternative on the t^{th} click; $c_{tnj} = 0$ otherwise. Let \vec{x}_{tnj} denote the characteristics for click-alternative j faced by consumer n on the t^{th} click. Let \vec{X}_{tn} be the set of \vec{x}_{tnj} 's up to an including the t^{th} click for all $j = 1$ to J_{tn} . Let \tilde{u}_{tnj} be the utility that consumer n obtains from clicking on the j^{th} click alternative on the t^{th} click. Let $\vec{\omega}_r$ be a vector of click-alternative-characteristic preferences for the r^{th} consumer segment and $\tilde{\epsilon}_{tnj}$ be an extreme value error such that $\tilde{u}_{tnj} = \vec{x}'_{tnj}\vec{\omega}_r + \tilde{\epsilon}_{tnj}$. Let Ω be the matrix of the $\vec{\omega}_r$'s. Let $\delta_{mn} = 1$ if the n^{th} consumer makes a purchase after seeing morph m ; $\delta_{mn} = 0$ otherwise.

The likelihood that the n^{th} respondent chooses clicks $\vec{c}_{T_n n}$ given the consumer belongs to segment r is given by:

$$(1) \quad \Pr(\vec{c}_{T_n n} | r_n = r, \Omega, \vec{X}_{T_n}) = \Pr(\vec{c}_{T_n n} | r_n = r) = \prod_{t=1}^{T_n} \prod_{j=1}^{J_{tn}} \left(\frac{\exp[\vec{x}'_{tnj}\vec{\omega}_r]}{\sum_{l=1}^{J_{tn}} \exp[\vec{x}'_{tnl}\vec{\omega}_r]} \right)^{c_{tnj}}$$

We estimate Ω from the calibration study by forming the likelihood over all respondents and by using standard maximum-likelihood methods or Bayesian methods. Denote these estimates by $\hat{\Omega}$.

In Morphing 1.0, we observe the consumer's clickstream up to the τ_o^{th} click. The unconditional prior probabilities, $\Pr_0(r_n = r)$ are observed in the calibration study or from website experience. Bayes Theorem provides:

$$(2) \quad q_{rn\tau_o}(\vec{c}_{\tau_o n}, \hat{\Omega}, \vec{X}_{\tau_o n}) \\ \equiv \Pr(r_n = r | \vec{c}_{\tau_o n}, \hat{\Omega}, \vec{X}_{\tau_o n}) = \frac{\Pr\{\vec{c}_{\tau_o n} | r_n = r, \hat{\Omega}, \vec{X}_{\tau_o n}\} \Pr_0(r_n = r)}{\sum_{s=1}^R \Pr\{\vec{c}_{\tau_o n} | r_n = s, \hat{\Omega}, \vec{X}_{\tau_o n}\} \Pr_0(r_n = s)}$$

For ease of exposition, we temporarily add the r subscript to δ_{rmn} to indicate a situation in which the segment, r , is known. Let p_{rmn} be the probability that consumer n in segment r , who experienced morph m , will make a purchase (or other success criterion). This probability is distributed: $f_n(p_{rmn}|\alpha_{rmn}, \beta_{rmn}) \sim p_{rmn}^{\alpha_{rmn}-1} (1 - p_{rmn})^{\beta_{rmn}-1}$ where α_{rmn} and β_{rmn} are parameters of the beta distribution. Updating implies $\alpha_{r,m,n+1} = \alpha_{rmn} + \delta_{rmn}$ and $\beta_{r,m,n+1} = \beta_{rmn} + (1 - \delta_{rmn})$. Normalizing the value of a purchase to 1.0, the expected immediate reward is $E[p_{rmn}|\alpha_{rmn}, \beta_{rmn}] = \alpha_{rmn}/(\alpha_{rmn} + \beta_{rmn})$.

Let G_{rmn} be the Gittins' index for the m^{th} morph for consumers in segment r , let $a \leq 1$ be the discount rate from one consumer to the next, and let $V_{Gittins}(\alpha_{rmn}, \beta_{rmn}, a)$ be the value of continuing with parameters a , α_{rmn} , and β_{rmn} . We table G_{rmn} by iteratively solving the Bellman equation.

$$(3) \quad V_{Gittins}(\alpha_{rmn}, \beta_{rmn}, a) = \max \left\{ \begin{aligned} & \frac{G_{rmn}}{1-a}, \frac{\alpha_{rmn}}{\alpha_{rmn} + \beta_{rmn}} [1 + aV_{Gittins}(\alpha_{rmn} + 1, \beta_{rmn}, a)] \\ & + \frac{\beta_{rmn}}{\alpha_{rmn} + \beta_{rmn}} aV_{Gittins}(\alpha_{rmn}, \beta_{rmn} + 1, a) \end{aligned} \right\}$$

When consumer segments are latent, we replace the Gittins' index with the expected Gittins' index, EGI_{mn} .

$$(4) \quad EGI_{mn} = \sum_{r=1}^R q_{rn\tau_o}(\vec{c}_{\tau_o n}, \widehat{\Omega}, \vec{X}_{\tau_o n}) G_{rmn}(\alpha_{rmn}, \beta_{rmn}, a)$$

For latent segments, the updating equations are based on “fractional observations.” Details are available in Hauser, Liberali, and Urban (2014).

$$(5) \quad \begin{aligned} \alpha_{r,m,n+1} &= \alpha_{rmn} + q_{rnT_n}(\vec{c}_{T_n n}, \widehat{\Omega}, \vec{X}_{T_n n}) \delta_{mn} \\ \beta_{r,m,n+1} &= \beta_{rmn} + q_{rnT_n}(\vec{c}_{T_n n}, \widehat{\Omega}, \vec{X}_{T_n n}) (1 - \delta_{mn}) \end{aligned}$$

For the Morphing 2.0 extension, let w_t be the weight for observation period t and let γ be the multiplicative switching cost. We add a t subscript to morphs such that m_{tn} indicates the morph seen by consumer n in the t^{th} observation period. To keep track of

morph changes, we define an indicator variable such that $\Delta_{m't_n tn} = 1$ if we change to morph m_{t_n}' for consumer n in period t ; $\Delta_{m't_n tn} = 0$ otherwise. Because the consumer may see many morphs, we drop the m subscript from δ_{mn} such that $\delta_n = 1$ if the consumer makes a purchase; $\delta_n = 0$ otherwise.

To determine when to morph, we solve a Bellman equation by backward recursion for each consumer. The immediate reward is the γ -discounted, weighted expected Gittins' index. The expectation uses $q_{rn}(\vec{c}_{t-1,n}, \widehat{\Omega}, \vec{X}_{t-1,n})$ because this inferred probability represents our expectations over all future clicks. The segment-conditional continuation reward is $V_t(m_{tn}^*, m_{t-1,n}, \vec{c}_{t-1,n}, \widehat{\Omega}, \vec{X}_{t-1,n} | r_{true,n} = s)$. It is computed by keeping track of morph changes for $\tau \geq t$. We take the expectation with respect to the probability of observing each consumer segment to obtain the unconditional reward. Let ψ_t be the probability of exit after the t^{th} observation period and let $\bar{\Psi}(S|t-1) = E_n[\prod_{s=t}^S (1 - \psi_s)]$, Then the Bellman equation is:

$$(6) \quad \begin{aligned} & V_t(m_{tn}^*, m_{t-1,n}, \vec{c}_{t-1,n}, \widehat{\Omega}, \vec{X}_{t-1,n}) \\ &= \max_{m_{tn}} \left\{ \begin{aligned} & \gamma^{\Delta_{m_{tn} tn}} w_t \sum_r q_{rn}(\vec{c}_{t-1,n}, \widehat{\Omega}, \vec{X}_{t-1,n}) G_{rm_{tn}n} \bar{\Psi}(t|t-1) + \\ & \sum_s [q_{sn}(\vec{c}_{t-1,n}, \widehat{\Omega}, \vec{X}_{t-1,n}) V_{t+1}(m_{t+1,n}^*, m_{tn}, \vec{c}_{t-1,n}, \widehat{\Omega}, \vec{X}_{t-1,n}, r.e. | s)] \bar{\Psi}(t+1|t) \end{aligned} \right\} \end{aligned}$$

We let $\eta_{mnt} = 1$ if consumer n saw morph m during the t^{th} observation period; $\eta_{mnt} = 0$ otherwise. Generalized fractional-observation updating becomes:

$$(7) \quad \begin{aligned} \alpha_{rm,n+1} &= \alpha_{rmn} + q_r(\vec{c}_{T_n n}, \widehat{\Omega}, \vec{X}_{T_n n}) \gamma^{N_{T_n}} \left(\sum_{t=1}^{T_n} \eta_{mnt} w_t \right) \delta_n \\ \beta_{rm,n+1} &= \beta_{rmn} + q_r(\vec{c}_{T_n n}, \widehat{\Omega}, \vec{X}_{T_n n}) \gamma^{N_{T_n}} \left(\sum_{t=1}^{T_n} \eta_{mnt} w_t \right) (1 - \delta_n) \end{aligned}$$