MANAGEMENT SCIENCE
Articles in Advance, pp. 1–23
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)

http://dx.doi.org/10.1287/mnsc.2014.1961
© 2014 INFORMS

Website Morphing 2.0: Switching Costs, Partial Exposure, Random Exit, and When to Morph

John R. Hauser
MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, hauser@mit.edu

Guilherme Liberali
Erasmus School of Economics, Erasmus University Rotterdam, 3000 DR Rotterdam, The Netherlands; and MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, liberali@mit.edu

Glen L. Urban
MIT Center for Digital Business, MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, glurban@mit.edu

Website morphing infers latent customer segments from clickstreams and then changes websites' look and feel to maximize revenue. The established algorithm infers latent segments from a preset number of clicks and then selects the best "morph" using expected Gittins indices. Switching costs, potential website exit, and all clicks prior to morphing are ignored. We model switching costs, potential website exit, and the (potentially differential) impact of all clicks to determine when to morph for each customer. Morphing earlier means more customer clicks are based on the optimal morph; morphing later reveals more about the customer's latent segment. We couple this within-customer optimization to between-customer expected Gittins index optimization to determine which website "look and feel" to give to each customer at each click. We evaluate the improved algorithm with synthetic data and with a proof-of-feasibility application to Japanese bank card loans. The proposed algorithm generalizes the established algorithm, is feasible in real time, performs substantially better when tuning parameters are identified from calibration data, and is reasonably robust to misspecification.

Data, as supplemental material, are available at http://dx.doi.org/10.1287/mnsc.2014.1961.

Keywords: automated marketing; Bayesian methods; clickstream analysis; dynamic programming; Internet marketing; optimization; switching costs; website design; website morphing
History: Received August 12, 2012; accepted March 19, 2014, by Eric Bradlow, special issue on business analytics. Published online in *Articles in Advance*.

1. Website Morphing and Its Limitations

Website morphing customizes the look and feel of a website to customers' cognitive styles. The basic algorithm, developed by Hauser, Urban, Liberali, and Braun (2009, hereafter HULB), combines Bayesian inference of latent cognitive-style segments with dynamic programming optimization to match website designs to customers. Bayesian inference on a customer's clickstream infers probabilities that the customer belongs to the latent segments. Using these probabilities and data from past purchases, the dynamic program automatically selects the best look and feel for the website for each customer. The "morph" assignment is (near) optimal in the sense that it balances learning about the best assignment for a segment with the profit that can be obtained by exploiting current knowledge of morph-to-segment purchase probabilities. HULB use data from a "calibration study" to simulate what would had happened had the BT Group implemented morphing on its broadband-sales website. HULB estimate that morphing would have increased revenue by $80 million.

Subsequently, Urban et al. (2014) adapted morphing to banner advertising. The only modification in the HULB algorithm was to account for multiple customer visits to the same website. In a field test with over 100,000 customers viewing over 450,000 banners on a CNET website, they report that banner morphing almost doubled click-through rates relative to a random assignment of banners. They also conducted a laboratory test on an automotive information and recommendation website to test the basic concept of morph-to-segment matching. The experiment replaced the automated algorithm with direct measurement in a four-to-five week longitudinal study. Click-through

rates on banners, as well as consideration and preference for Chevrolet-branded vehicles, increased significantly when morphs were matched to customer segments. The Chevrolet application expanded the definitions of customer segments to include the stage of the automotive buying process (collecting data, comparing vehicles, committing to a purchase).

The HULB algorithm, used in the BT Group and CNET tests, assumes that only the customer's final clicks on the website affect the probability of purchase (or, in the case of banners, click-through rates). The algorithm assumes further that the act of switching the look and feel of a website or banner has no effect on the probability of purchase. And, finally, the algorithm does not take into account that customers might exit the website, sometimes after a relatively few clicks.

This paper proposes an improved algorithm that accounts for three phenomena: the impact of all clicks on a customer's experience, switching costs, and random exit. We do so by transforming an inherently hard switching-cost problem, with a curse-of-dimensionality computational explosion, into a formulation that accounts for the phenomena and can run between clicks on a real website. We demonstrate that accounting for these phenomena can increase performance dramatically. For example, a counterfactual policy simulation, based on a proof-of-feasibility application to a Japanese bank website, suggests sales would have increased relative to random morph assignment. The increase was 63% of that obtainable by perfect information and over five times the 12% achievable by the HULB algorithm. (Systematic simulations, reanalyzing the HULB application, suggest average improvements of about 1.8× that of HULB.)

The improved algorithm nests the HULB algorithm in the sense that it reduces to the HULB algorithm for specific (and we feel extreme) values of key parameters. When parameters can be identified from a calibration study, the algorithm does at least as well as the HULB algorithm. For many reasonable parameter values, the improved algorithm outperforms the HULB algorithm substantially. The improved algorithm almost always outperforms the HULB algorithm even if parameter values are chosen incorrectly.

We begin by reviewing the HULB algorithm, then present the behavioral model, and propose the improved algorithm. Subsequent sections describe the synthetic-data analyses and policy simulations based on a Japanese bank application. We close with a discussion of unsolved problems.

2. Brief Review of the HULB Website Morphing Algorithm

This section presents an abridged version of the HULB website morphing algorithm. To stay consistent, we adopt the notation used by both HULB and Urban et al. (2014). Appendix A summarizes all notation. Greater details on derivations are available in HULB. A user's guide, code, pseudocode, example questionnaires, and data are available in the supplemental material to this paper (available at http://dx.doi.org/10.1287/mnsc.2014.1961).

2.1. The Constructs of Morphing: Morphs, Segments, Clickstream, Calibration

We first provide a qualitative overview of the constructs used in morphing. Detailed examples are available in both HULB and Urban et al. (2014) and their online appendices.

2.1.1. Morphs. Morphs refer to alternative implementations of the overall look and feel of a marketing instrument such as a website (HULB) or banners (Urban et al. 2014). For simplicity of notation, this paper focuses on the look and feel of websites. Creating alternative morphs imposes a fixed cost on website or banner design—Urban et al. (2014) report an estimated cost of $250,000 for the many Chevrolet banners tested. This fixed cost (and sample size issues discussed later) means that most applications use a moderate number of morphs. HULB use eight.

2.1.2. Customer Segments. Customers are classified into a set of mutually exclusive and collectively exhaustive segments under the hypothesis that different segments respond differently to the morphs. The BT Group, CNET, and Chevrolet applications classified customers by cognitive styles; in the Chevrolet application customers were also classified by the stage of their buying process. Typically, in a calibration study prior to implementing the morphing algorithm, segment membership can be observed with intrusive measurement. During day-to-day implementation on a website, (latent) segments cannot be observed directly and must be inferred (probabilistically) from the customer's clickstream. As with morphs, practical considerations require that the number of segments be moderate—16 for the BT Group application, 4 for the CNET application, and 12 for the Chevrolet application. The selection of segments is a modeling decision as discussed extensively in the market segmentation literature.

2.1.3. Clickstream. As a customer explores a website in day-to-day operations, we observe the customer's clickstream. In particular, on any webpage, the customer must choose among a number of links. Some links are pictures, some are text, some take the customer to another webpage, some open a comparison tool, etc. We call these *click alternatives*, and we observe the customer's sequential choice of click alternatives. So that we might use the customer's sequential choices to infer the customer's segment probabilities, we code each click alternative by a series of click characteristics.

For example, HULB code each click alternative by 11 characteristics, some of which are based on evaluation by independent judges (e.g., graphs or text), some of which code specific targets (e.g., an analytic tool), and some of which code actions (e.g., post a comment).

2.1.4. Outcomes. Our goal is to maximize an outcome such as sales (BT Group) or click-through (CNET). The algorithm works with other binary outcomes such as aware or not aware, prefer or not prefer, intend to purchase or not intend to purchase, or a Boolean combination of such measures. The outcome is observed at the end of the customer's visit to the website. For example, a BT Group customer might either purchase broadband service or leave the website. Typically, we code a customer as leaving the website if there is no activity in 30 minutes. The BT Group application considered only one visit to the website, but the CNET application used cookies to track multiple visits. In CNET, a successful outcome occurred if the customer clicked through on at least one of the visits. Algorithms seek to assign morphs to customers by maximizing the time discounted value of successful outcomes from the current customer and all future customers.

2.1.5. Calibration Study. Prior to implementing the morphing algorithm in day-to-day operations, some parameters must be estimated. To obtain data to estimate these parameters, we recruit a sample of customers to answer questions before and after visiting the website. During the calibration study, morphs are assigned randomly to respondents. These calibration-study questions enable us to assign calibration customers to segments and estimate a model that, conditioned on segment membership, predicts customers' click-alternative choices as a function of the characteristics of the click alternatives. Because it is not feasible to ask such calibration-study questions for day-to-day visitors, we use a Bayesian model to estimate day-to-day customers' latent segment probabilities. We refer to customers in the calibration study as respondents to avoid confusion.

2.2. Posterior Probabilities of (Latent) Segment Membership

2.2.1. Notation. Let n index customers, r index segments, m index morphs, and t index clicks for each customer. Capital letters R, M, and T_n denote totals. (We do not need N for customers. We use N later for a different construct.) Let c_{tn} denote the tth click by nth customer, and let $\vec{c}_{tn} = \{c_{1n}, c_{2n}, \ldots, c_{tn}\}$ denote the vector of clicks up to and including the tth click. At each click choice, the customer faces J_{tn} click alternatives as denoted by c_{tnj}, where j indexes click alternatives. We let $c_{tnj} = 1$ if customer n clicks the jth click alternative on the tth click and $c_{tnj} = 0$ otherwise. Let $\vec{x}_{tnj}$ denote the characteristics for click alternative j faced by customer n on the tth click. Let $\vec{X}_{tn}$ be the set of the $\vec{x}_{tnj}$ up to an including the tth click for all $j = 1$ to J_{tn}. Let $\tilde{u}_{tnj}$ be the utility that customer n obtains from clicking on the jth click alternative on the tth click. Let $\vec{\omega}_r$ be a vector of click-alternative-characteristic preferences for the rth customer segment, and let $\tilde{\epsilon}_{tnj}$ be an extreme value error such that $\tilde{u}_{tnj} = \vec{x}'_{tnh}\vec{\omega}_r + \tilde{\epsilon}_{tnj}$. Let Ω be the matrix of the $\vec{\omega}_r$'s.

2.2.2. Calibration of Preferences for Click-Alternative Characteristics. In the calibration study, we observe the customer's segment, r, the customer's click-alternative choices, $\vec{c}_{T_n n}$, and the click-alternative characteristics the customer faced at each click-alternative choice, $\vec{x}_{tnj}$. The extreme value error gives us the standard logit model with the following likelihood for the nth respondent in the calibration study:

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

The likelihood over all respondents in the calibration study is simply the product of the individual-respondent likelihoods. With Equation (1), we obtain estimates (maximum-likelihood methods) or Bayesian posteriors (Markov chain Monte Carlo sampling, with priors) of Ω. (HULB use Bayesian posteriors; Urban et al. 2014 use maximum-likelihood methods.) We denote the estimates (or the mean of the posterior distributions, if Bayesian methods are used) by $\hat{\Omega}$. Because morph assignments are made in real time, it is not feasible to sample from the full posterior distribution when Bayesian methods are used.

2.2.3. Probabilities That a Customer Belongs to Each (Latent) Segment. In day-to-day website operations, we treat $\hat{\Omega}$ as known. We observe the customer's clickstream, $\vec{c}_{\tau_o n}$, up to the τ_oth click, and we observe the relevant click-alternative characteristics, the $\vec{X}_{\tau_o n}$. We seek to estimate the probability that the nth customer belongs to segment r for all $r = 1$ to R. Denote these probabilities, after the τ_oth click, by $q_{rn\tau_o}$. From the calibration study we know the unconditional prior probabilities, $\Pr_0(r_n = r)$, that the nth customer belongs to segment r. (Some websites might recalibrate these priors periodically.) Given the observed clickstream, clickstream characteristics, and $\hat{\Omega}$, we use Equation (1) to obtain $\Pr(\vec{c}_{\tau_o n} \mid r_n = r, \hat{\Omega}, \vec{X}_{tn})$. Bayes' theorem provides the following:

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

Equation (2) (and the imbedded Equation (1)) runs sufficiently fast on most computers that we obtain the $q_{rn\tau_o}(\vec{c}_{\tau_o n}, \hat{\Omega}, \vec{X}_{\tau_o n})$ between clicks on the website.

2.3. Choosing the Optimal Morph If We Knew the Customer's Segment

Customer n's true segment is latent, but we find it easier to explain the HULB algorithm for a hypothetical case that assumes we know customer n's true segment. We relax this assumption in §2.4. Let p_{rmn} be the probability that customer n in segment r, who experienced morph m, will make a purchase (or other success criterion). We represent our knowledge, prior to customer 1, about p_{rm1} with a beta distribution that has parameters α_{rm1} and β_{rm1}. Specifically, $f_1(p_{rm1} \mid \alpha_{rm1}, \beta_{rm1}) \sim p_{rm1}^{\alpha_{rm1}-1}(1-p_{rm1})^{\beta_{rm1}-1}$. Define p_{rmn}, α_{rmn}, and β_{rmn} accordingly. After each customer, we update to the posterior distribution, $f_n(p_{rmn} \mid \alpha_{rmn}, \beta_{rmn}, \delta_{rmn})$, where α_{rmn} and β_{rmn} are sufficient to summarize data from customers 1 to $n-1$. Because α_{rmn} and β_{rmn} are used to make decisions about customer n, they do not yet include observations for customer n.

Let $\delta_{rmn} = 1$ if the nth customer is in segment r and makes a purchase after seeing morph m; $\delta_{rmn} = 0$ otherwise. (We temporarily add the r subscript to δ_{rmn} to indicate that the segment, r, is known.) Then, because the binomial outcomes are naturally conjugate to beta priors, Appendix B shows that $\alpha_{rm,\,n+1} = \alpha_{rmn} + \delta_{rmn}$ and $\beta_{rm,\,n+1} = \beta_{rmn} + (1-\delta_{rmn})$. Normalizing the value of a purchase to 1.0 implies that the immediate reward we expect from the nth customer is $E[p_{rmn} \mid \alpha_{rmn}, \beta_{rmn}] = \alpha_{rmn}/(\alpha_{rmn} + \beta_{rmn})$.

Assigning the optimal morph to the nth customer is more complicated than simply maximizing the immediate reward. Whenever we assign a morph m to a customer in segment r and observe an outcome, we update the posterior distribution for p_{rmn}. The updated distribution enables us to make better decisions in the future. The dynamic decision problem must balance immediate rewards with the knowledge gained that enables better decisions in the future.

HULB demonstrate that this dynamic decision problem, for known r_n, is the classic multiarmed bandit problem studied by Gittins (1979). Gittins (1979) proved that the problem is "indexable" and that the optimal solution is to compute an index for each arm and choose the arm with the largest index. (Indexability has a technical definition, but in lay terms, an "arm" of a multiarmed bandit is indexable if, when an index increases, the set of states for which choosing the arm is optimal does not decrease. Indexability implies an index strategy is feasible, although it does not guarantee an index strategy is optimal. Very often an index strategy is near optimal.) Because the detailed derivations are available in Gittins (1979), summarized in HULB, and otherwise widely known, we do not repeat them here. We repeat only the basic equations.

Let G_{rmn} be the Gittins index for the mth morph assuming the customer is in segment r, and we have updated α_{rmn} and β_{rmn} based on all customer purchases up to but not including the nth customer. If $a \le 1$ is the discount rate from one customer to the next and if $V_{\text{Gittins}}(\alpha_{rmn}, \beta_{rmn}, a)$ is the value of continuing with parameters a, α_{rmn}, and β_{rmn}, then G_{rmn} solves the following Bellman equation:

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

Intuitively, the right-hand side chooses the maximum of a fixed "arm" and an uncertain "arm." The fixed "arm" pays the expected value Gittins index over all future n. The uncertain arm pays a unit reward if the outcome is a success and nothing if it is not. A success occurs with an expected probability of $\alpha_{rmn}/(\alpha_{rmn} + \beta_{rmn})$. The uncertain arm also pays the value of continuing to explore optimally for future n, but discounted by a. The continuation values account for updating the beta distribution.

Equation (3) does not have an analytic solution, but we readily compute the Gittins indices with a simple iterative numeric algorithm. (We reuse code developed by Gittins. Code is available in Gittins et al. 2011.) For a given a, we table G_{rmn} as a function of α_{rmn} and β_{rmn}. If the segment r were known, the HULB algorithm would simply look up the Gittins indices for all m and choose the morph with the largest index.[1] This strategy would lead to the maximum long-term profit taking both exploration and exploitation into account. In the special case where we are limited to using the same website or banner for everyone ($R = 1$), we can use the Gittins index algorithm to choose the optimal website or banner for each n.

2.4. Choosing a Morph When Segments Are Partially Observable

2.4.1. Partially Observable Markov Decision Process. When a customer's segment is not known with certainty, the dynamic program changes. Because the latent segments is only partially observable, the optimization problem requires we solve a partially observable Markov decision process (POMDP). The state space is Markov because the full history is summarized by the α_{rmn}'s, the β_{rmn}'s, and the latent segments.

[1] As n gets large, $G_{rmn} \to \alpha_{rmn}/(\alpha_{rmn} + \beta_{rmn})$. We table α_{rmn} and β_{rmn} up to the point where this difference is small, that is, to a maximum of 3,000. Because the function is smooth, we use linear interpolation for noninteger values. See example in Appendix 3 of HULB.

2.4.2. Expected Gittins Index. Krishnamurthy and Mickova (1999) prove that the POMDP is an indexable decision process. Krishnamurthy and Mickova (1999) establish further that if we compute the expectation of the Gittins index over our uncertainty about the customer's latent segment and choose the morph that has the largest expected Gittins index (EGI), then, in many cases, the EGI policy will be very close to optimality. More critically, because the EGI is easy to compute, we can select a (near) optimal morph in real time (between clicks on the website). HULB provide simulation evidence that the EGI strategy works extremely well for website morphing.

Specifically, the EGI algorithm replaces the Gittins index with the expected Gittins index, EG_{mn}, and chooses the morph with the largest EG_{mn}, where

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

The HULB algorithm requires that we choose arbitrarily the click on which to morph. HULB use $\tau_o = 10$, and Urban et al. (2014) use $\tau_o = 5$. Equation (4), with minor modification, also enables us to choose the best starting morph for each customer. We simply replace $q_{rn\tau_o}(\vec{c}_{\tau_o n}, \hat{\Omega}, \vec{X}_{\tau_o n})$ with q_{rno}, our beliefs about the customer's segment prior to observing any clicks. HULB use $q_{rno} = \Pr_0(r_n = r)$, but it would also be feasible to use periodic updates.

2.4.3. Updating Beliefs When Segments Are Partially Observable. Updating beliefs when the customer's segment is known is particularly easy because the binomial distribution and the beta distribution are naturally conjugate. When the segments are latent, Appendix B demonstrates that updating is no longer naturally conjugate. However, we can still update if we consider "fractional observations"; that is, if we observe a success, δ_{mn}, conditioned on the customer having seen morph m, we consider this as a fractional success for each latent customer segment, r. (Note that the outcome, δ_{mn}, is not conditioned on r because the segment is latent.) The fractional success is $q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})\delta_{mn}$ for each $r = 1$ *to* R. The binomial distribution is well defined for fractional observations and naturally conjugate to the beta distribution. HULB update via

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

Updating occurs when the customer leaves the website; thus we use all clicks, $\vec{c}_{T_n n}$. Appendix B provides complete derivations as well as arguments that these updating formulae lead to posterior distributions that converge to a mass point at the true values of the $p_{rm,\,\text{true}}$'s as $n \to \infty$. HULB provide simulation evidence that the fractional updating formulae lead to effective and profitable morph-to-segment assignments. These simple formulae enable the morphing algorithm to run in real time between a customer's clicks on the website. With future, faster computers, one might improve the updating formulae.

2.5. Experience with the HULB Algorithm

Website morphing imposes a high initial fixed cost. The estimated $80 million in incremental revenue for the BT Group, the observed banner click-through lifts of 80%–100% for CNET's context-matched banners, and the observed 30% lift in Chevrolet brand consideration justify the initial fixed cost for these high-traffic websites. Morphing also requires high traffic because of its convergence properties.

A sale by the BT Group to the $(n+1)$st customer is worth almost as much as a sale to the nth customer. The HULB example with 100,000 visitors per annum implies a discount rate of $a = 0.999999$ (HULB, p. 208). When a is close to 1.0, the optimal strategy includes substantial exploration. With known segments, HULB report that their algorithm explored heavily for the first 1,000 customers; it did not stop exploring until the 3,000th customer. They illustrate convergence with their Figure 3 (HULB, p. 209). With 16 equal-sized segments, the HULB algorithm would likely stabilize around the 50,000th customer.

Success probabilities for banners are typically much lower than success probabilities for sales on a morphing website—the order of 0.003 in Urban et al. (2014) versus 0.38 in HULB. As a result, Gittins indices take longer to stabilize in banner morphing than in website morphing. Nonetheless, the Urban et al. (2014) experience is informative. In their application with 100,000 customers, the algorithm stabilized for the largest segments, but was still exploring for the smallest segment of approximately 9,000 customers.

The experience by HULB and Urban et al. (2014) suggest that we should evaluate performance up to approximately 10,000 customers per segment. Thus, the four-segment synthetic data experiments in this paper examine performance from 1 to 40,000 customers per simulation—10,000 customers per segment.

3. Behavioral Theory Improvements: Key Assumptions

The HULB algorithm has no way to determine the click on which to morph. Instead, both HULB and Urban et al. (2014) choose an arbitrary click. This limitation would only be optimal if the morph seen by customers prior to the τ_oth click had negligible impact on successful outcomes (sales or click-through). The implications of this implicit assumption were never tested.

Because Gittins' (1979) solution requires that switching costs are zero (Banks and Sundaram 1994), and because bandit problems with switching costs are often computationally difficult, HULB assumed that customers experienced no switching costs when the website morphed. Indeed, without special structure, switching costs cause the optimization problem to become NP complete (Jun 2004, p. 526; Ny and Feron 2006, Theorem 1).

To address switching costs and the decision of when to morph, we specify a theory of customer behavior. To date, no such theory exists with respect to morphing. We seek a theory that is parsimonious, captures the essence of the phenomena, and is sufficiently "conjugate" that a when-to-morph (WTM) algorithm is feasible in real time between clicks on the website. To the extent that these assumptions can be improved, our tests are conservative. Future elaborations might improve website morphing further.

3.1. Assumption 1: Switching Costs

An extensive literature in psychology and marketing studies the cognitive cost of switching. Papers in psychology, beginning with Jersild (1927), generally study the cognitive cost of completing a task (e.g., Spector and Biederman 1976). Meiran (2000) establishes that switching among methods of response in a computer task imposes cognitive loads on respondents. In marketing, switching costs are well established in many sales contexts (e.g., Weiss and Anderson 1992), appear to apply for cognitive costs (Jones et al. 2000, 2002), and affect purchase intention. For websites, Balabanis et al. (2006, p. 217) suggest that "cognitive search costs in online shopping environments can be significant," and Johnson et al. (2003, p. 63) report that "perceived switching costs . . . create a cognitive 'lock-in'" for websites. It is likely that the cognitive costs of switching morphs reduce the benefit of matching a website's look and feel to a customer segment.

The cost of switching is salient with a hypothetical example. Suppose that the initial morph for the nth customer is graphical and focused. The nth customer experiences this morph and learns to search based on a graphical-focused look and feel. Now suppose that subsequent clicks by the customer suggest instead that the best morph is verbal and general. If we were to switch morphs, the customer might become confused and have to relearn his or her search strategy. Although the verbal-general morph might have been best for the customer had the website had those characteristics from the beginning, it may not be best after the customer has learned to use a graphical-focused website. In general, it is such path dependence that makes switching-cost optimization problems NP complete.

Additive switching cost are common in the literature, and algorithms exist (e.g., Banks and Sundaram 1994, Dushochet and Hongler 2006, Jun 2004), but additive switching costs impose path dependence and make it infeasible to determine when to morph in real time. On the other hand, a multiplicative switching cost can be factored out in a Bellman equation. As we demonstrate below, we can solve a problem with multiplicative switching costs in real time. Specifically, we assume that a switch in the look and feel of the website lowers the customer's purchase probability. The switch lowers the purchase probability by a factor of γ, where $\gamma \leq 1$. In this formulation, the HULB algorithm is a special case that assumes $\gamma = 1$.

Fortunately, multiplicative switching costs have descriptive advantages. First, because predicted outcomes, p_{rmn}'s, are bounded between 0 and 1, a multiplicative factor does not violate that bound, whereas an additive factor might. Second, although untested, we expect that the amount by which a low probability is lowered by a switching cost will be less than the amount by which a high probability is lowered by a switching cost. For example, suppose a switch lowers p_{rmn} from 0.900 to 0.810. A comparable proportional cost would lower p_{rmn} from 0.090 to 0.081, whereas a comparable additive cost would lower p_{rmn} from 0.090 to 0.000. Likely the former is more realistic.

One might generalize γ to assume that $\gamma_{mm'}$ is a function of the morphs from which and to which the customer switches. We leave that generalization to future research because (1) it would introduce severe state dependence that would likely make the dynamic program infeasible, (2) it would introduce a substantial measurement burden requiring a large number of parameters (56 parameters in the BT Group application and 132 parameters in the Chevrolet application), and (3) the HULB algorithm treats morphs as independent. Nonetheless, morph-specific switching costs are an interesting and challenging extension.

3.2. Assumption 2: The Impact of Being Exposed to Multiple Morphs

The HULB algorithm assumes that only the last morph affects purchase probabilities. Implicitly, morphs prior to the switch have zero impact on purchase probabilities. However, suppose that the system morphs after the 10th click and the customer makes a purchase decision after the 20th click. There is no reason to assume that the first 10 clicks have less impact than the last 10 clicks. The last 10 clicks may have more impact (recency), less impact (primacy), or equal impact. To allow various assumptions, we assign weights, w_t, to clicks to account for the differential impact of early versus late morphs. Let $\vec{w}$ be the vector of these weights. (Section 3.5 discusses other models.) Setting $w_t = w$ for all t implies equal impact. Setting w_t equal to an increasing (decreasing) function of t assigns greater impact to later (earlier) morphs. The HULB algorithm

assumes a special case for $\vec{w}$. In particular, the HULB algorithm implicitly sets w_t equal to zero for $t \le \tau_o$ and $w_t = w$ for $t > \tau_o$.

In equation form, if customer n sees morph m_{tn} at the tth click (or the tth observation period), we generalize the HULB assumption to

$$p_{rn} = \sum_t w_t p_{rm_{tn}}. \quad (6)$$

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). Switching costs, when relevant, apply for all t in Equation (6).

3.3. Assumption 3: Variation in the Number of Clicks for Each Customer's Visit

Some customers find the information they need quickly, make a purchase decision, and exit the website. Other customers visit many areas of the website, gather extensive information, and leave later. We cannot assume, nor do the data support, an assumption that all customers stay for all observation periods. The number of clicks does not seem to be particularly correlated with either the morph given or the purchase probability, but rather reflects the (unobservable) information needs of the customer.

Let ψ_{tn} be the probability that customer n leaves on the tth click given that customer n has already made $t-1$ clicks. For example, before we observe the first click, we expect the customer to leave after that click with probability ψ_{1n}, to leave after the second click with probability $(1-\psi_{1n})\psi_{2n}$, to leave after the third click with probability $(1-\psi_{1n})(1-\psi_{2n})\psi_{3n}$, and so on. For each possibility, we normalize the effective impact weights, w'_t, to account for the number of clicks before exit. This model generalizes the implicit assumptions of HULB that $\psi_{tn} = 0$ for $t < T_n$ and $\psi_{T_n,n} = 1$ for some $T_n \gg \tau_o$.

The only datum we observe for each customer is the click at which he or she leaves; thus we need a reasonably parsimonious model that balances nonstationarity over t with heterogeneity over n. In §6.6 we demonstrate how to estimate the parameters of the model from the calibration study. Recall that morphs are assigned randomly during the calibration study. We obtain maximum-likelihood estimates based on the observed times of exit.

3.3.1. Model 1: Heterogeneity over Customers. Model 1 assumes that ψ_{tn} is beta distributed over customers with parameters α_ψ and β_ψ but independent of t: $f_\psi(\psi_{tn} \mid \alpha_\psi, \beta_\psi) \sim \psi_{tn}^{\alpha_\psi - 1}(1-\psi_{tn})^{\beta_\psi - 1}$. Given this assumption, we can readily calculate the probability that a randomly chosen customer will leave after any given click. We obtain the predicted probability of leaving by integrating out the heterogeneity. For example, simple calculus provides the probability of leaving after the first click as $\alpha_\psi/(\alpha_\psi + \beta_\psi)$, the probability of leaving after the second click as $\alpha_\psi\beta_\psi/[(\alpha_\psi + \beta_\psi + 1)(\alpha_\psi + \beta_\psi)]$, after the third click as $\alpha_\psi\beta_\psi(\beta_\psi + 1)/[(\alpha_\psi + \beta_\psi + 2)(\alpha_\psi + \beta_\psi + 1)(\alpha_\psi + \beta_\psi)]$, etc. Using these values, we calculate the probability that a customer remains on the website through the tth click.

3.3.2. Model 2: Homogeneity over Customers, But Nonstationarity over Clicks. Assume $\psi_{tn} = \psi_t$ for all n.

3.3.3. Model 3: Hybrid Model. Respondents might remain on the website for some initial clicks, but, after those initial clicks, revert to a different ψ_{tn} that is heterogeneous in n (as in Model 1).

3.3.4. Model 4: Extensions. We might assume that exit probabilities depend on the customer's segment, the terminal morph, or the morph history, e.g., ψ_{trm}, $\psi_{trmm'}$, etc. For example, $\psi_{trmm'}$ might be a function of the dissimilarity between morph m and morph m'. These extensions complicate the optimization challenge and are left to future research.

3.4. Selecting Values of the Tuning Parameters

The switching discount (γ) and the impact weights ($\vec{w}$) are tuning parameters in the more generalized algorithm. They must be selected *before* the algorithm is used to morph a website (in day-to-day operations). The time of exit is easy to observe in the calibration study. We estimate the parameters of a model for exit probabilities (ψ_{tn}'s) using standard methods. (Exit seems more tied to the customer's information needs than the morph given; hence calibration data are likely sufficient for these parameters.)

The tuning parameters, γ and $\vec{w}$, require either managerial judgment or experiments during the calibration study. In a calibration study, segment membership is measured directly, therefore the true r_n is known among calibration respondents. Applications to date vary morphs randomly over customers to estimate the priors, p_{rm1}. A more complex experimental design might vary morphs randomly over customers and randomly switch morphs within customers. With a sufficient sample size, estimates of the tuning parameters, $\hat{\gamma}$, $\hat{\vec{w}}$, and the $\hat{p}_{rm1}$'s, are all identified from observed customer outcomes (sales or click-throughs). For example, if there are eight morphs, four segments, and four observation periods, we need to estimate 37 parameters (32 $\hat{p}_{rm1}$'s, 4 $\hat{w}_t$'s, and $\hat{\gamma}$). For each segment, some customers see one morph for all periods. Some customers see two morphs with switches varied over $t = 1$, 2, *or* 3. A full factorial would be 8×4 (single morph combinations) plus $8 \times 7 \times 4 \times 3$ (two morph combinations) for a total of 704 combinations—more than enough to identify

37 parameters. If a feasible algorithm could be developed to account for different $\gamma_{mm'}$'s, they, too, would be identified, albeit with a large sample size requirement.

We might also update the tuning parameters after initial experience on the website, say, after observing 50,000 customers. Although the optimal morph allocations are based on the assumed values of the tuning parameters, we might still identify $\hat{\gamma}$ and $\vec{\hat{w}}$ based on the likelihood principle. See related discussions in Hauser and Toubia (2005) and Liu et al. (2007). Basically, in vivo, a morphing algorithm will have assigned various morphs to consumers at various time periods and will have done so based on known rules from known data. We observe δ_{mn} for every consumer who was so assigned. If the likelihood principle is assumed, it should be feasible to estimate γ and $\vec{w}$. Future algorithms might also explore optimal experimentation to learn the tuning parameters while learning about the p_{rmn}'s.

3.5. Alternative Assumptions

Making the decision of when (whether, how often) to morph is a challenging optimization problem. We must solve the problem optimally or near optimally, and we must do so in real time between customer clicks. Modeling switching costs, impact weights, and exit probabilities with separable functions enables rapid solutions. We believe our assumptions are simple, realistic, flexible, and well matched to the data available in website-morphing applications. However, other assumptions are possible. For example, we might make γ an increasing function of the clicks before a switch and a decreasing function of the clicks after a switch, or we might allow γ to depend on the number of prior switches. Equation (6) might be replaced by a nonlinear function. Either γ or $\vec{w}$ might be modeled as heterogeneous over segments or over customers.

All website morphing applications to date used clicks rather than clock time because clock time adds unobserved variance due to network speed, distractions while browsing, and variations in reaction time. Although clicks are correlated with clock time in our field experiment ($\rho = 0.57$, $p < 0.01$), future applications might model γ, $\vec{w}$, and ψ_{tn} as functions of clock time rather than clicks. With experience and new research, we might find one or more of these extensions feasible and profitable. For now, we believe our assumptions are reasonable, robust, and generalize website morphing.

4. When to Morph: Explicit Endogenous Decisions on Timing

The HULB algorithm identifies the best morph to give to each customer based on a fixed number of customer clicks. We can do better with an algorithm that determines endogenously the number of clicks to observe before morphing. The improved WTM algorithm also allows, but does not require, more than one change in morphs. In this section, t indexes observation periods that may be more than one click.

4.1. Dynamic Decision Problem

Figure 1 illustrates the WTM decision problem for the case where the customer makes a purchase (or leaves the website) after four observation periods. (The theory applies when the number of observation periods is a random variable; four morphs is just an illustration.) Specifically, during observation period t_n, the website displays morph m_{tn}. The respondent makes clicks, c_{tn}, while exploring the website, and we update our beliefs about the customer's segment, $q_{rn}(\vec{c}_{tn}, \hat{\Omega}, \vec{X}_{tn})$. Using the new information, and anticipating more information from subsequent decision periods, we decide which morph, $m_{t+1,n}$, to display in the next decision period. To keep track of morph changes, we define $\Delta_{m'_{tn}t}$ as an indicator variable such that $\Delta_{m'_{tn}tn} = 1$ if we *change* to morph m'_{tn} for customer n in period t. With this notation, the total number of morph changes for customer n is $N_{tn} = N_{t-1,n} + \sum_{m'_{tn}} \Delta_{m'_{tn}tn}$. We represent the purchase decision by $\delta_n = 1$ if the customer makes a purchase and $\delta_n = 0$ if the customer does not make a purchase. We dropped the m subscript from HULB's δ_{mn} to allow for the fact that more than a single morph may have affected outcomes for customer n.

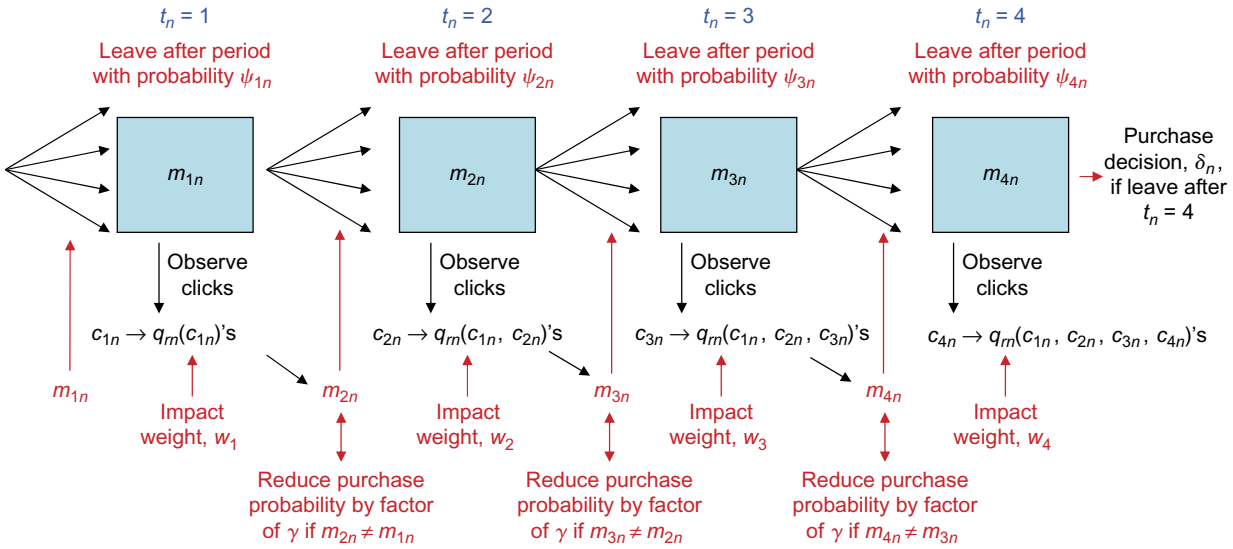
Figure 1 illustrates the basic dilemma. The longer we wait to morph, the more clicks we observe. More clicks enable us to identify better customer n's segment and, hence, the best morph for customer n. However, when the impact weights imply that early morph experience affects purchase probabilities or if the customer is likely to leave early, we want to get to that best morph as rapidly as feasible. The problem is compounded because switches are costly.

4.2. Formulating a Feasible Bellman Equation

Optimizing morph decisions for the problem in Figure 1 is challenging. It is even more challenging when embedded within a dynamic program to learn the best morph to assign for each customer segment (as in HULB). For example, Asawa and Teneketzis (1996, p. 329) caution: "inclusion of a switching penalty drastically changes the nature of the bandit problem . . . . The optimal policy is not given by an index rule anymore." Although heuristics exist using multiple indices (e.g., Dusonchet and Hongler 2006), the number of indices explodes exponentially in a morphing problem were there are, potentially, tens of thousands of costly switches (one or more switches for each customer).

The problem is further compounded when Bayesian updating is used to identify customer segments. Each switch presents new click opportunities, which customers choose probabilistically based on their (latent) customer segment and the morph decision. We might anticipate how our decisions in observation period t

Figure 1 When-to-Morph Decision Problem (Example Customer, Four-Period Illustration)



Notes. Choose morph (m_{tn}) at each decision period to maximize sales, δ_n. The q_{rn}'s are also a function of $\hat{\Omega}$ and $\vec{X}_{tn}$. Some arguments are suppressed in this figure to avoid complexity.

affect the observations that update $q_{rn}(\vec{c}_{\tau n}, \hat{\Omega}, \vec{X}_{\tau n})$ for $\tau > t$. This problem quickly becomes intractable. We need to finesse the explosion in the number of potential click paths. Finally, any solution must take into account that customers leave stochastically between observation periods.

We address these challenges with an algorithm that runs between a customer's clicks on the website. Our solution recognizes that the two decisions, morph-to-segment assignment (HULB) and when to morph, are based on learning and optimization that occur over vastly different time scales. The HULB optimization is based on learning from observed outcomes of tens of thousands of customers, whereas the WTM optimization problem is based on learning customer segments from clicks *within* a customer's visit. Because Gittins indices, G_{rmn}, summarize the rewards from optimal assignment in all future periods (and are only updated when a customer exits), the indices remain constant between clicks and serve to summarize the value of assigning morph m to customer n. This heuristic strategy is analogous (but not identical) to optimal solutions in the branching bandit literature. (We expand these intuitive arguments in §4.2.5.)

To make the WTM algorithm feasible between clicks, we exploit the likelihood principle. Specifically, after the tth observation period, the best Bayesian estimate of customer n's purchase probabilities after all observation periods ($\vec{c}_{T_n n}$) is based only on clicks up to and including the tth observation period ($\vec{c}_{tn}$). This principle enables us to finesse the explosion in click opportunities that is dependent on our decisions for customer n after observation period t. To formulate the Bellman equation, we consider carefully what we know about customer n's segment, when we know it, and how this affects the future.

4.2.1. Immediate Reward. The first simplification comes from linear impact weights, which enable us to separate outcomes by observation period. The second simplification comes from the multiplicative nature of switching costs, which enables us to factor them out. The third simplification comes when we recognize that, independently of future clicks, our best estimate at t of the terminal probabilities, $q_{rn}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})$, is $q_{rn}(\vec{c}_{t-1,n}, \hat{\Omega}, \vec{X}_{t-1,n})$. In other words, the $q_{rn}(\vec{c}_{t-1,n}, \hat{\Omega}, \vec{X}_{t-1,n})$ represent our expectations over all future clicks. With these simplifications, we write the expected immediate reward, $EIR(m_{tn}, m_{t-1}, \vec{c}_{t-1,n}, \hat{\Omega}, \vec{X}_{t-1,n})$, as

$$\begin{aligned} &EIR(m_{tn}, m_{t-1,n}, \vec{c}_{t-1,n}, \hat{\Omega}, \vec{X}_{t-1,n}) \\ &\quad = \gamma^{\Delta_{m_{tn}}} w_{tn} \sum_r E_{c_{tn}, c_{t+1,n}, \ldots, c_{Tn} \mid \vec{c}_{t-1,n}} q_{rn}(\vec{c}_{Tn}, \hat{\Omega}, \vec{X}_{T_n n}) G_{rm_{tn}} \\ &\quad = \gamma^{\Delta_{m_{tn}}} w_t \sum_r q_{rn}(\vec{c}_{t-1,n}, \hat{\Omega}, \vec{X}_{t-1,n}) G_{rm_{tn}}. \end{aligned}$$

4.2.2. Value of Continuing Optimally. To formulate the value of continuing optimally, we recognize that the evolution of $q_{rn}(\vec{c}_{\tau n}, \hat{\Omega}, \vec{X}_{\tau n})$ depends on the true customer segment. However, at period t (and even at period T_n), we are uncertain about the true segment. We have an expectation over $r_{\text{true},n}$ for the purchase probabilities, and we have an expectation over $r_{\text{true},n}$ for the evolution of $q_{rn}(\vec{c}_{\tau n}, \hat{\Omega}, \vec{X}_{\tau n})$ for $\tau \geq t$.

To keep track of these expectations, we temporarily write the evolution of the clicks as conditioned on $r_{\text{true},n}$.

If $m^*_{t+1,\,n}(r_{\text{true}})$ is the morph we would choose at $t+1$ if the true segment for customer n was $r_{\text{true},\,n}$, then the continuation value is

$$V_{t+1}\big(m^*_{t+1,\,n}(r_{\text{true},\,n}),\, m_{tn},\, \vec{c}_{t-1,\,n},\, \hat{\Omega},\, \vec{X}_{t-1,\,n} \mid r_{\text{true},\,n}\big) \equiv \max_{m_{t+1,\,n}} E_{c_{tn},\,c_{t+1,\,n},\ldots,c_{Tn}\,|\,\vec{c}_{t-1,\,n}} V_{t+1}(m_{t+1,\,n},\, m_{tn},\, \vec{c}_{t-1,\,n},\, c_{tn},\, c_{t+1,\,n},\ldots,\, c_{T_n n},\, \hat{\Omega},\, \vec{X}_{T_n n} \mid r_{\text{true},\,n}).$$

But $r_{\text{true},\,n}$ is not known when we are making the decision at t, so we account for this dependence for $\tau \geq t$ by taking an expectation over this unknown variable at t; that is, we take an expectation over the true segments when we compute the continuation value for $\tau \geq t$. Going forward to any observation period, τ, under the assumption that the true segment is $r_{\text{true},\,n}$, we would like to anticipate how the $q_{rn}(\vec{c}_{\tau n},\, \hat{\Omega},\, \vec{X}_{\tau n})$ will evolve for any future sequence of morphs. However, in practical situations, the state space is so large that we cannot anticipate the future clickstream for $\tau \geq t$. At the start of the tth decision period, our best estimate of customer n's segment probabilities over all future clickstream paths remains $q_{rn}(\vec{c}_{t-1,\,n},\, \hat{\Omega},\, \vec{X}_{t-1,\,n})$. In other words, V_τ depends on $\vec{c}_{t-1,\,n}$ for $\tau \geq t$. Backward induction does the rest.

Putting it all together enables us to optimize the current period's morph, m_{tn}, by backward induction using a computationally rapid Bellman equation. Naturally, we keep track of the $\Delta_{m_{\tau n}\tau n}$'s for $\tau \geq t$ when we compute the conditional values, $V_\tau(m^*_{\tau n},\, m_{\tau-1,\,n},\, \vec{c}_{\tau-1,\,n},\, \hat{\Omega},\, \vec{X}_{\tau-1,\,n} \mid r_{\text{true},\,n} = s)$:

$$\begin{aligned} V_t(m^*_{tn}, m_{t-1,\,n}, \vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n}) = \max_{m_{tn}} \Bigg\{ & \gamma^{\Delta_{m_{tn}}} w_t \sum_r q_{rn}(\vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n}) G_{r m_{tn} n} \\ & + \sum_s \big[q_{sn}(\vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n}) \\ & \cdot V_{t+1}(m^*_{t+1,\,n}, m_{tn}, \vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n} \mid r_{\text{true},\,n} = s) \big] \Bigg\}. \end{aligned} \tag{7}$$

4.2.3. Variation in the Number of Periods for Each Customer's Visit. Purely for ease of exposition, we derived the WTM Bellman equation assuming a fixed number of observation periods. It is relatively simple to generalize the Bellman equation for variation in the number of periods for each customer's visit. The assumption of §3.3 provides the separability necessary. Let $\Psi_n(S \mid t-1)$ be the probability that customer n is still at the website at observation period S given that customer n was at the website at observation period $t-1$. The expectation over customers is $\bar{\Psi}(S \mid t-1)$ because we do not know customer n's exit probability prior to observing an exit. Using any of the random-exit models in §3.3, we calculate this probability via $\bar{\Psi}(S \mid t-1) = E_n[\prod_{s=t}^{S}(1-\psi_{ns})]$. The random exit ($r.e.$) Bellman equation is then the following equation, where the impact weights are renormalized to reflect the periods in which the customer remains on the website:

$$\begin{aligned} V_t(m^*_{tn}, m_{t-1,\,n}, \vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n} \mid r.e.) = \max_{m_{tn}} \Bigg\{ & \gamma^{\Delta_{m_{tn}}} w_t \sum_r q_{rn}(\vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n}) G_{r m_{tn} n} \bar{\Psi}(t \mid t-1) \\ & + \sum_s \big[q_{sn}(\vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n}) \\ & \cdot V_{t+1}(m^*_{t+1,\,n}, m_{tn}, \vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n}, r.e. \mid s) \big] \\ & \cdot \bar{\Psi}(t+1 \mid t) \Bigg\}. \end{aligned} \tag{8}$$

4.2.4. Modified Bayesian Updating. The HULB algorithm assumes that only the last morph seen by customer n affects the probability of a successful outcome for customer n. For real-time computation, HULB treated the observed outcome, δ_{mn}, as a series of fractional outcomes as determined by the (latent) segment probabilities, $q_{rn}(\vec{c}_{T_n n},\, \hat{\Omega},\, \vec{X}_{T_n n})$. The WTM algorithm is based on a more general behavioral assumption that allows impact weights for all observation periods. We therefore generalize the updating procedure (Equation (5), §2.4.3). As demonstrated in Appendix B, we can generalize HULB and treat each combination of a period and a latent segment as a fractional observation.

Let $\eta_{mnt} = 1$ if customer n saw morph m during the tth observation period, and let $\eta_{mnt} = 0$ otherwise. The fractional observation for period t, morph m, and segment r is based on the η_{mnt}, the impact weights w_t, the switching-cost penalty, and $q_r(\vec{c}_{T_n n},\, \hat{\Omega},\, \vec{X}_{T_n n})$. Appendix B motivates the following (practical) updating equations:

$$\begin{aligned} \alpha_{rm,\,n+1} &= \alpha_{rmn} + q_r(\vec{c}_{T_n n},\, \hat{\Omega},\, \vec{X}_{T_n n}) \gamma^{N_{T_n}} \cdot \left(\sum_{t=1}^{T_n} \eta_{mnt} w_t \right) \delta_n, \\ \beta_{rm,\,n+1} &= \beta_{rmn} + q_r(\vec{c}_{T_n n},\, \hat{\Omega},\, \vec{X}_{T_n n}) \gamma^{N_{T_n}} \cdot \left(\sum_{t=1}^{T_n} \eta_{mnt} w_t \right) (1 - \delta_n). \end{aligned} \tag{9}$$

4.2.5. Motivation for Expected Gittins Index Decoupling. The WTM optimization (Equations (7) and (8)) is feasible because we decoupled the decision on when to morph for a customer from optimal experimentation between customers. We cannot prove that decoupling retains the near optimality of the expected Gittins index, but we can examine the synthetic-data experiments and the Japanese bank counterfactual. The WTM algorithm does quite well relative to the HULB algorithm and relative to an upper bound of

perfect information. The WTM algorithm achieves a net present value that is 60%+ of that achieved by perfect information. No algorithm could do better than perfect information.

There are intuitive reasons to believe that decoupling will be near optimal. First, the observed δ_n cannot affect the decision of when to morph for customer n. The Gittins indices, G_{rmn}'s, and the distributions of the purchase probabilities, p_{rmn}'s, are updated *after* customer n leaves the website. Second, the WTM decision is based on learning about a customer's segment during a visit, whereas the morph-to-segment decision (HULB) is based on learning about the purchase probabilities using the terminal segment-membership probabilities. Third, learning through Gittins indices happens many orders of magnitudes slower than decisions of when to morph. Empirical evidence (§2.5) suggests that Gittins indices stabilize after 40,000–100,000 customers depending on the number of morphs and segments. On the other hand, decisions about when to morph depend on clickstream observations on the order of fractions of a single customer's website visit. Fourth, decoupling is analogous (but not equivalent) to the "branching bandits" literature, where researchers have proven that it is often optimal to replace the uncertain outcome of an indexable decision process with its Gittins index (Bertsimas and Niño-Mora 1996, Tsitsiklis 1994). For example, Weber (1992) analyzes a "super process" in which an initial bandit process has rewards that are themselves bandit processes. To solve the superprocess dynamic program optimally, Weber (1992) replaces the outcome of each secondary bandit process with its Gittins index. He then uses those indices as rewards when computing the indices for the arms in the primary bandit process.

At minimum, we expect the WTM algorithm will provide substantial improvements relative to the HULB algorithm. We examine the magnitude of such improvements in §5.

4.3. Summary of the Algorithmic Improvements

The HULB algorithm has proven successful on calibration-study-based synthetic data representing the BT Group website and for a field experiment morphing banners on CNET. However, the HULB algorithm assumes that only the last morph matters, that there are no switching costs, and that customers do not leave the website randomly. Our proposed WTM algorithm generalizes the HULB algorithm to address those issues. The algorithm is based on a series of behavioral assumptions that we feel are reasonable and capture the relevant phenomena. The WTM algorithm runs in real time between clicks on a website and identifies when to morph while retaining the ability to allocate the best morph to each customer. The WTM algorithm requires "tuning" parameters for impact weights and switching costs. The tuning parameters can be estimated with a sufficiently large sample in the calibration study. Lacking a calibration study, the parameters can be set by managerial judgment. We now examine the WTM algorithm with synthetic data and with a counterfactual policy simulation based on a proof-of-feasibility application.

5. Synthetic Data Experiments

If customers are described accurately by the behavioral assumptions, the principle of optimality suggests that the WTM algorithm will outperform the HULB nested algorithm. We use synthetic data experiments to examine the amount of the improvement. We also examine robustness by comparing the HULB algorithm to the WTM algorithm when the WTM algorithm uses incorrect values of the tuning parameters.

5.1. Reanalysis of the HULB BT Group Simulations

HULB tested their algorithm with synthetic data chosen to mimic behavior on the BT Group website. In particular, they used the $\hat{p}_{rm}$ estimated from the calibration data to create $p_{rm,\,\text{true}}$'s. If morph m was assigned to a customer in segment r after the first τ_o clicks, HULB drew a binomial outcome based on $p_{rm,\,\text{true}}$ and updated accordingly. To simulate Bayesian segment identification, HULB drew customer click-alternative choices with Equation (1) using the $\hat{\Omega}$ estimated from the calibration data (up to 20 webpages and 16 click alternatives per webpage). The $q_{rn\tau_o}(\vec{c}_{\tau_o n}, \hat{\Omega}, \vec{X}_{\tau_o n})$ were based on the simulated clickstream and Equation (2).

We begin by reexamining the HULB BT Group simulations. We examine a hypothetical world in which there are four observation periods of five webpages each (20 webpages as in the HULB simulations). Data are generated with $\gamma = 0.95$ and $\vec{w} = (0.25, 0.25, 0.25, 0.25)$. We compare the proposed WTM algorithm with the HULB algorithm, which implicitly assumes $\gamma = 1$ and $\vec{w} = (0, 0, 0, 1)$. This example illustrates that there exists at least one case where the WTM algorithm substantially improves revenue. Because the calculated rewards include switching costs, the rewards in Table 1 are smaller than those reported by HULB.

We select four customer segments and simulate 10,000 customers from each segment (a total of 40,000 synthetic customers per simulation). Visual inspection of the G_{rmn} plots indicates that 10,000 synthetic customers per segment are sufficient to observe performance for finite n and at (or near) convergence. Although the variance over simulations is small, we simulate each algorithm 10 times and take the average.

Table 1 summarizes the relative improvements. We report the net present value of the rewards using the discount factor from HULB, $a = 0.999999$. We also report the rewards over the last of 40,000 synthetic customers as an indication of the rewards after the G_{rmn}'s have

Table 1 Reanalysis of BT Group Website Morphing (with Switching Costs and Impact Weights)

	Reward at convergence[a]	Percent improvement in reward[b] (%)	Net present value (NPV)[c]	Percent improvement in NPV[d] (%)
Baseline for no morphing	0.3171	0.0	0.3108	0.0
Website morphing				
Website morphing (HULB)	0.3392	33.4	0.3306	30.6
Website morphing with improved algorithm	0.3567	60.0	0.3442	51.6
Perfect information ($n \to \infty$)[e]	0.3831	100.0	0.3754	100.0

Notes. For these comparisons, switching costs are $\gamma = 0.95$. Impact weights are $\vec{w} = (0.25, 0.25, 0.25, 0.25)$.

[a]Rewards for the last 400 customers out of 40,000 customers.

[b](Reward due to algorithm − reward at no morphing)/(reward due to perfect information − reward at no morphing).

[c]Net present value of the rewards divided by the number of periods.

[d](NPV due to algorithm − NPV at no morphing)/(NPV due to perfect information − NPV at no morphing).

[e]This is the upper bound. Applications do not have perfect information on either customer segments or purchase probabilities.

converged to the p_{rmn}'s. The relative improvement in the net present value of rewards is substantial (51.6% for WTM versus 30.6% for HULB). The relative improvement is even larger at convergence (60.0% for WTM versus 33.4% for HULB).

HULB (p. 212) report that a 1% increase in sales on the BT Group website is worth approximately \$4 million in revenue. Based on this valuation, the improvement due to the HULB algorithm over random allocation is approximately \$25.5 million. The improvement of the WTM over the HULB algorithm is worth approximately \$17.5 million.

5.2. Systematic Comparisons Using a Full-Factorial Design for Tuning Parameters

We examine whether the WTM algorithm improves outcomes for a wide range of switching costs and impact weights. We vary switching costs systematically from $\gamma = 0.80$ to $\gamma = 1$. We choose $\gamma = 0.80$ as a lower value because HULB report an improvement of approximately 20%. Such an advantage would not be sufficient if $\gamma < 0.80$. The upper value corresponds to the HULB implicit assumption. We vary impact weights systematically from equally valuing all observation periods ($w_t = w \ \forall\, t$) to favoring only the last observation period ($w_t = 0 \ \forall\, t < T_n$, $w_{T_n} = 1$). Specifically, we vary $\vec{w}$ over the following values: (0.25, 0.25, 0.25, 0.25), (0.10, 0.15, 0.25, 0.50), (0.02, 0.05, 0.18, 0.75), and (0, 0, 0, 1). We plot performance for $n = 1$ to 40,000 at increments of 400 to illustrate both small-sample properties and properties near the convergence of the G_{rmn}'s. We do not test sensitivity to random exit (ψ) because the tuning parameters for random exit are easy to infer from the calibration data (as in §6.6). The discount rate, a, is determined directly from the firm's discount rate and the number of website visitors per year.

Table 2 and Figure 2 compare the net present value of projected revenues and the values of revenue at convergence, for the WTM algorithm compared with the HULB algorithm. We report comparisons for all 20 (5×4) combinations of γ and $\vec{w}$. Projected revenues are averaged over 10 replications—a total of 400 simulations and 16 million synthetic customers.

In Figure 2, the vertical axes are the projected revenue per customer with the revenue for a success normalized to 1.0. The horizontal axes are the number of synthetic customers reported every 400th customer. In all cases, both algorithms converge smoothly. By 40,000 synthetic customers, convergence is usually complete. (For some values of the tuning parameters, revenues are still increasing slowly at 40,000. It is unlikely we would gain further insight beyond 40,000 synthetic customers per experimental cell.) The case where the WTM algorithm matches the HULB algorithm, $\gamma = 1$ and $\vec{w} = (0, 0, 0, 1)$, is shown in the lower right. As predicted, the WTM and HULB algorithms perform identically. In all other cases, the WTM algorithm performs better than the HULB algorithm—sometimes substantially better. Improvements in the net present value average 1.78 times higher. They vary from 1.19 at $\gamma = 1$, $\vec{w} = (0.02, 0.05, 0.18, 0.75)$ to an almost fourfold increase of 3.89 at $\gamma = 0.80$, $\vec{w} = (0, 0, 0, 1)$. Increases at convergence are less dramatic, but still substantial. Ratios vary from 1.17 to 2.17 with an average of 1.59.

Table 2 provides insight on when the additional complexity of WTM has the largest impact. The relative increase is most sensitive to switching costs, with the largest increase when switching costs are largest. The effect of impact weights varies and interacts with switching costs. For modest switching costs, the relative increase is largest for equal impact weights. For large switching costs, the relative impact is largest when the impact weights favor the last period. Overall, the effect of impact weights is less than that of switching costs.

5.3. Robustness Tests When the WTM Tuning Parameters Differ from True Values

We now compare projected revenues when the WTM algorithm assumes incorrect values of the tuning parameters. Our synthetic data experiments are illustrative. A full-factorial crossing all misspecifications with all true values would require $(4 \times 5) \times (4 \times 5) = 400$ cases for a total of 2,000 simulations and 320 million synthetic respondents. Such a large number of simulations would

Table 2 Proposed When-to-Morph Algorithm vs. the HULB Algorithm as Switching Costs (γ) and Impact Weights ($\vec{w}$) Vary

	Ratio of improvement in net present value, WTM vs. HULB					Percent improvement in net present value due to the WTM algorithm (%)					Percent improvement in net present value due to the HULB algorithm (%)				
$\vec{w}$'s	$\gamma = 0.80$	$\gamma = 0.85$	$\gamma = 0.90$	$\gamma = 0.95$	$\gamma = 1$	$\gamma = 0.80$	$\gamma = 0.85$	$\gamma = 0.90$	$\gamma = 0.95$	$\gamma = 1$	$\gamma = 0.80$	$\gamma = 0.85$	$\gamma = 0.90$	$\gamma = 0.95$	$\gamma = 1$
(0.25, 0.25, 0.25, 0.25)	2.08	1.90	1.78	1.69	1.36	38	43	49	52	51	18	23	27	31	38
(0.10, 0.15, 0.25, 0.50)	3.03	2.67	2.02	1.66	1.28	44	51	56	62	63	14	19	28	37	50
(0.02, 0.05, 0.18, 0.75)	3.31	2.65	2.14	1.51	1.19	51	60	64	66	75	15	23	30	44	63
(0.00, 0.00, 0.00, 1.00)	3.89	2.48	1.82	1.36	1.00	62	64	68	73	77	16	26	37	54	77
	Ratio of improvement in value at $n = 40{,}000$, WTM vs. HULB					Percent improvement in value at $n = 40{,}000$ due to the WTM algorithm (%)					Percent improvement in value at $n = 40{,}000$ due to the HULB algorithm (%)				
$\vec{w}$'s	$\gamma = 0.80$	$\gamma = 0.85$	$\gamma = 0.90$	$\gamma = 0.95$	$\gamma = 1$	$\gamma = 0.80$	$\gamma = 0.85$	$\gamma = 0.90$	$\gamma = 0.95$	$\gamma = 1$	$\gamma = 0.80$	$\gamma = 0.85$	$\gamma = 0.90$	$\gamma = 0.95$	$\gamma = 1$
(0.25, 0.25, 0.25, 0.25)	1.74	1.77	1.78	1.79	1.49	41	47	54	60	62	23	27	30	33	41
(0.10, 0.15, 0.25, 0.50)	1.85	1.97	1.78	1.56	1.30	46	56	62	69	74	25	29	35	44	57
(0.02, 0.05, 0.18, 0.75)	1.93	1.94	1.80	1.45	1.17	55	66	71	73	84	29	34	39	52	72
(0.00, 0.00, 0.00, 1.00)	2.17	1.82	1.61	1.28	1.00	68	71	77	82	87	31	39	48	64	87

be hard to display and provide little additional insight. Instead we select 12 interesting cases and examine whether the WTM algorithm is reasonably robust to parameter misspecification. We provide our code so that other researchers might run synthetic data experiments for other true and/or assumed values of the tuning parameters (see the supplemental material). Our 10-replicate analyses are based on an additional 240 simulations and 9.6 million synthetic customers.

Figure 3 displays the results. The first five plots hold $\vec{w}$ constant, match $\vec{w}$ to the implicit assumptions of HULB, and vary γ. We set $\gamma_{\text{true}} = 0.95$. In all cases, the WTM algorithm outperforms the HULB algorithm. The upper left plot suggests that when $\gamma_{\text{true}} = 0.95$, it is better to misspecify the tuning parameter ($\gamma = 0.80$, WTM) than to assume there are no switching costs ($\gamma = 1$, HULB).

The next four plots hold γ constant, match γ to the implicit assumptions of HULB, and vary $\vec{w}$. We set $\vec{w}_{\text{true}} = (0.10, 0.15, 0.25, 0.50)$. The WTM algorithm outperforms the HULB algorithm. Finally, the last three plots favor the HULB algorithm. Whereas γ_{true} and the $\vec{w}_{\text{true}}$ match the HULB assumptions, γ and the $\vec{w}$ in WTM are misspecified. The HULB algorithm does slightly better, as anticipated, but not by much. The loss due to using a misspecified WTM algorithm is much smaller than the loss due to ignoring switching costs and impact weights.

The 12 plots are indicative. If γ_{true} and $\vec{w}_{\text{true}}$ differ from the values assumed in the WTM and HULB algorithms, then the WTM algorithm does substantially better than the HULB algorithm. If γ_{true} and $\vec{w}_{\text{true}}$ match the HULB algorithm, but γ and $\vec{w}$ are misspecified in the WTM algorithm, then the HULB algorithm does only slight better. These results suggest that the WTM algorithm is likely robust to misspecification. This is likely the case because the HULB algorithm assumes extreme values for the tuning parameters: the HULB algorithm assumes no switching costs ($\gamma = 1$) and ignores all but the last observation period ($w_t = 0$ for all but $t = T_n$). As is often the case in modeling, it is better to attempt to model phenomena, even if done so imprecisely, than it is to ignore the phenomena altogether (e.g., Little 1966, 1970).

Naturally, results might change for different $\hat{p}_{rm}$'s and other values of γ_{true} or $\hat{w}_{\text{true}}$. We expect WTM to perform particularly well when the website is designed so that we learn the customer's segment substantially faster than τ_o. The effect would be particularly strong when there are substantial differences by segment in the effects of the morphs on outcomes. We see one such case in the next section. The website was designed so that segments could be identified quickly.

Figure 2 Proposed Algorithm (WTM) vs. HULB as Switching Costs, γ, and Impact Weights, $\vec{w}$, Vary

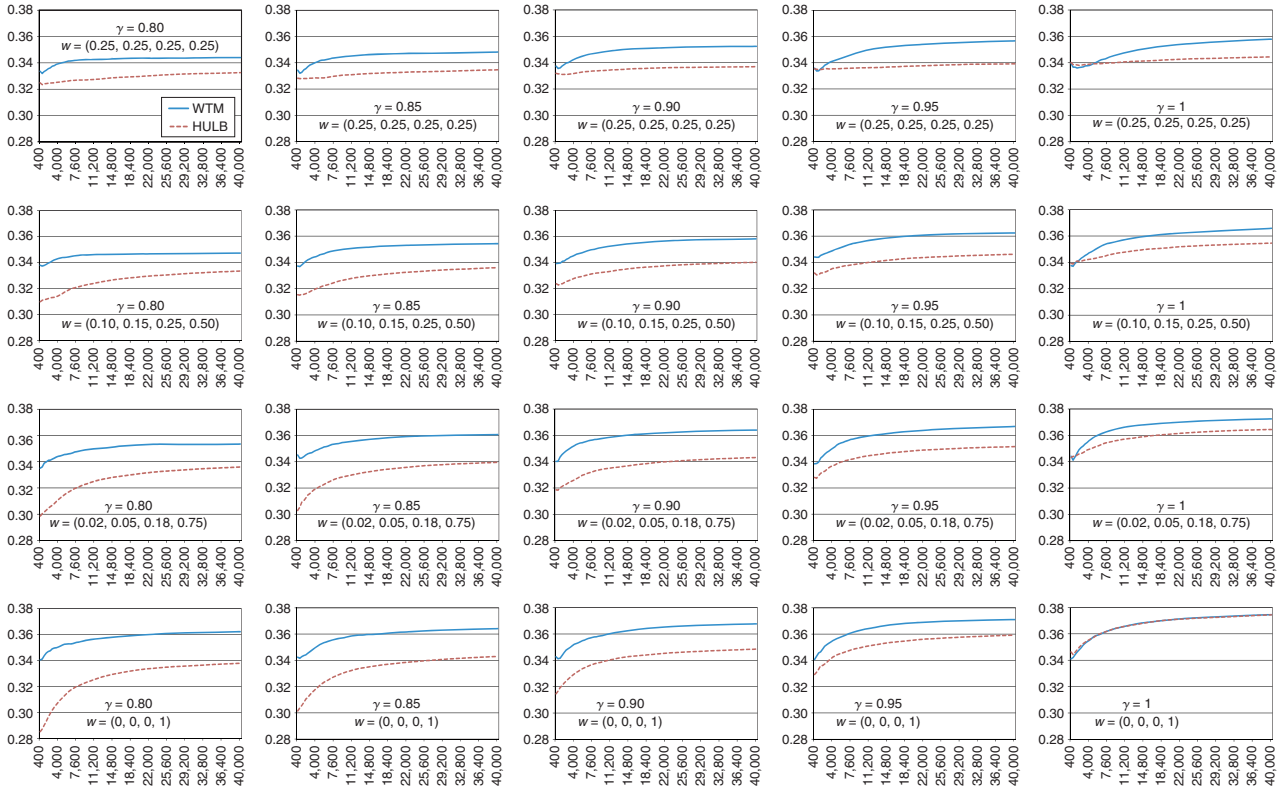
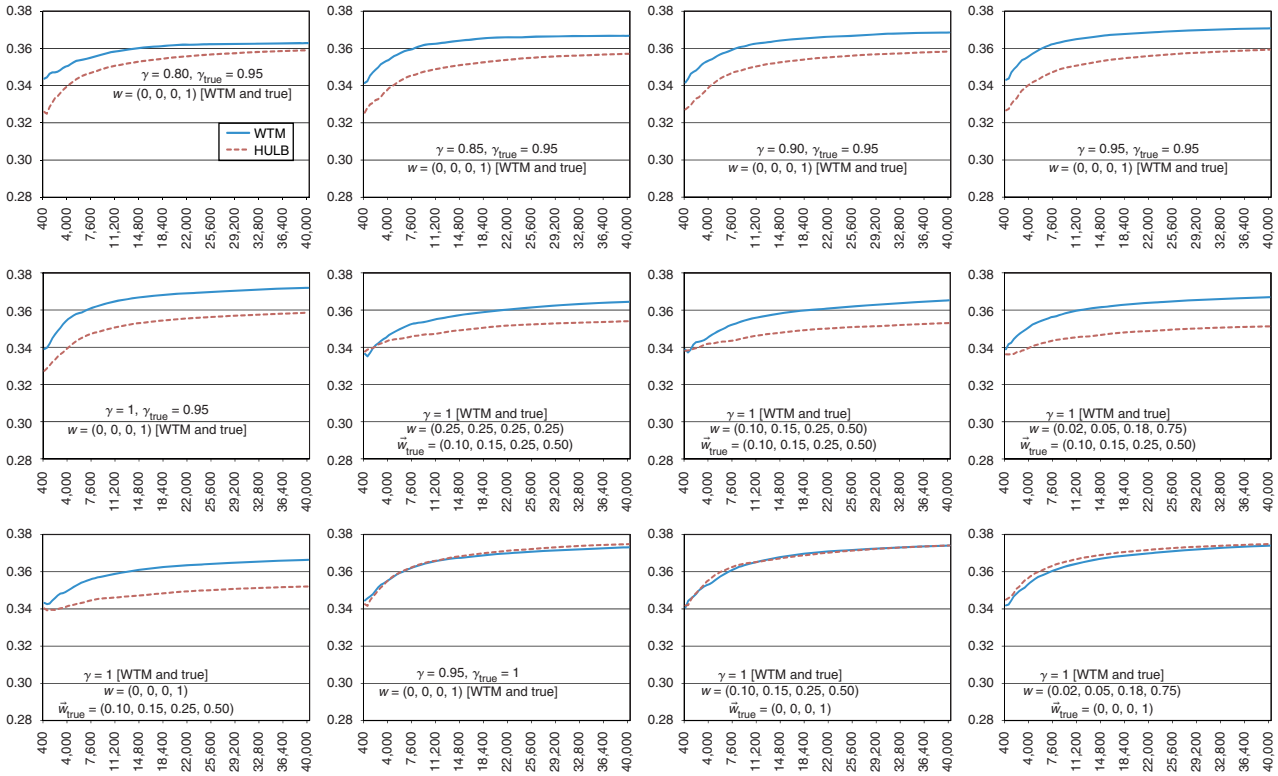


Figure 3 Representative Sensitivity Analyses



Note. This figure shows the results when true values of switching costs, γ, and impact weights, $\vec{w}$, do not match assumptions in the WTM algorithm.

6. Proof-of-Feasibility Implementation to a Japanese Bank Website

The WTM algorithm requires that we solve a dynamic program between clicks (or observation periods) while customers are actively using a website. In this section, we use data from a Japanese Bank website to examine whether evaluating when to morph is feasible in real time on a real website. We also examine whether the WTM-versus-HULB comparisons using BT Group data can be replicated using $\hat{p}_{rm}$ values from another website. This section also provides a roadmap to the practical implementation of morphing.

6.1. Context: Suruga Bank's Card-Loan Website

In Japan, customers prefer "card loans" rather than carrying a balance on their credit cards. (Japanese banks do not allow overdrafts.) The borrower receives a cash card with a balance of ¥3–¥5 million and pays interest when the funds are withdrawn. The terms of card loans vary among banks and are often confusing. Some banks offer low interest and high limits, but a more difficult screening process, whereas other banks offer higher interest and lower limits, but an easier screening process. In 2006–2007, the leading card loan banks were Orix, which spent ¥13.6 billion, mostly on banner advertising, and Acom, which spent ¥10.9 billion, mostly on television advertising ($\$1 \approx ¥95$ in that time period).

Suruga Bank is a Japanese commercial bank in the greater Tokyo area. Unlike most commercial banks, it focused on retail banking for more than 20 years. Suruga began a virtual bank in 1999, one of the first Japanese banks to do so. By 2008, its online presence had grown to 10 virtual branches and 8 virtual alliances (Tokoro 2008, p. 7). In that time period, Suruga was less well-known than other Japanese banks, spending approximately 1/10th that of Acom and Orix on advertising (¥1.4 million; Tokoro 2008, p. 17). As part of an overall strategy to reach more customers, Suruga developed a customer advocacy website on which it presented the best products from all competitors. By using a strategy of openness and honesty, Suruga sought to demonstrate that its products (low interest rates, high limits, but a careful screening process) would meet the needs of many customers. Suruga's managers found website morphing intriguing and authorized a small-scale field experiment.

6.2. Customer Segments and Estimated Click-Characteristic Preferences

We began with a calibration study to identify cognitive-style segments and estimate customers' preferences for click-alternative characteristics. In March 2008, 5,454 customers were drawn from a panel of customers maintained by Interface Asia. Of these, 3,340 were not interested in card loans or did not meet the age requirements. Those who qualified (2,114 respondents) were offered ¥200 and invited to visit an experimental website and complete a survey. Of these, 502 respondents (23.7%) completed the calibration survey, which included a requirement to browse the website for at least 2.5 minutes and for at least 10 clicks.

Table 3 Estimation of Click-Characteristic Preferences ($\hat{\Omega}$)

	Holistic vs. analytic	Impulsive vs. deliberative
Cognitive and cultural characteristics		
Pictures and graphs	0.88	−1.24
Technical, detailed content	−0.66	1.70
Textual content	1.10	−1.28
Options and alternatives	−0.13	1.96
Popular trends	0.11	−0.72
"You-directed" language	0.45	−3.61
Hierarchical images	0.24	−0.51
Functional characteristics		
Provide information	−0.53	1.56
Analytic tool	−0.60	−1.44
Graphical elements	−0.89	2.73
Website areas		
Advisor	−0.38	5.19
Fast solutions	−0.70	4.85
Learn information	0.11	1.89
Forum	0.37	2.48

Notes. Results are based on a maximum-likelihood logit analysis of Equation (1). $\chi^2_{42} = 2{,}505.5$ and $U^2 = 33.9\%$. Constants not shown.

By analyzing answers to the questions that were chosen to measure cognitive styles, we identified four segments as defined by two bipolar (ipsative) multi-item scales. The 2×2 categorization was impulsive versus deliberative and holistic versus analytic. We estimated the segment-specific click-characteristic preferences, $\hat{\Omega}$, from customers' observed clickstreams. Table 3 reports $\hat{\Omega}$ for the 14 click-alternative characteristics that were tracked. The cognitive and cultural characteristic values were based on ratings by six independent judges who were blind to the hypotheses of the research (reliability, 0.84). The functional characteristics and website areas were binary variables. The logit model (Equation (1)) was strongly significant ($p < 0.001$) and explained 33.9% of the uncertainty ($U^2 = 0.339$) based on 2,827 click observations. The estimated $\hat{\Omega}$ was intuitive. For example, impulsive visitors prefer links to fast solutions, advisors, and forums, but not analytic tools or "content directly addressed to you."

6.3. Website Design

Website designers (native Japanese speakers) developed four morphs that varied on the number of graphs, the amount of technical content, the amount of textual content, the number of options and alternatives presented, the amount of content on popular trends, the amount of "you-directed" content, formal versus informal Japanese language, and hierarchical versus

Figure 4 Example Screens from the Suruga Bank Website



egalitarian images. Although the designers did their best to develop morphs that would be more effective for specific segments, the best morph for each segment is determined automatically by the morphing algorithm. The website designers also developed click alternatives that would help identify segment membership. For example, on the opening page (Figure 4, first panel), customers enter the site by choosing one of two pictures. The pictures are designed to appeal to different customer segments. In another example, customers choose from six ways to obtain information (Figure 4, second panel) with the hope that different segments will choose different paths. To avoid an obvious demand artifact, the website was not identified as a Suruga Bank website. We expected that these design features would enable the Bayesian engine to identify customer segments more rapidly than was possible with the BT Group website.

6.4. Outcome Measure and Tuning Parameters

Because we were unable to sell card loans in the experiment, the algorithm maximized customers' requests for more information. In particular, $\delta_n = 1$ if the customer clicked on "send me more information." The tuning parameters were set conservatively by managerial judgment ($\hat{\gamma} = 0.99$, $\hat{w}_t = w \ \forall t$).

6.5. Sample Used in the Suruga Proof-of-Feasibility Application

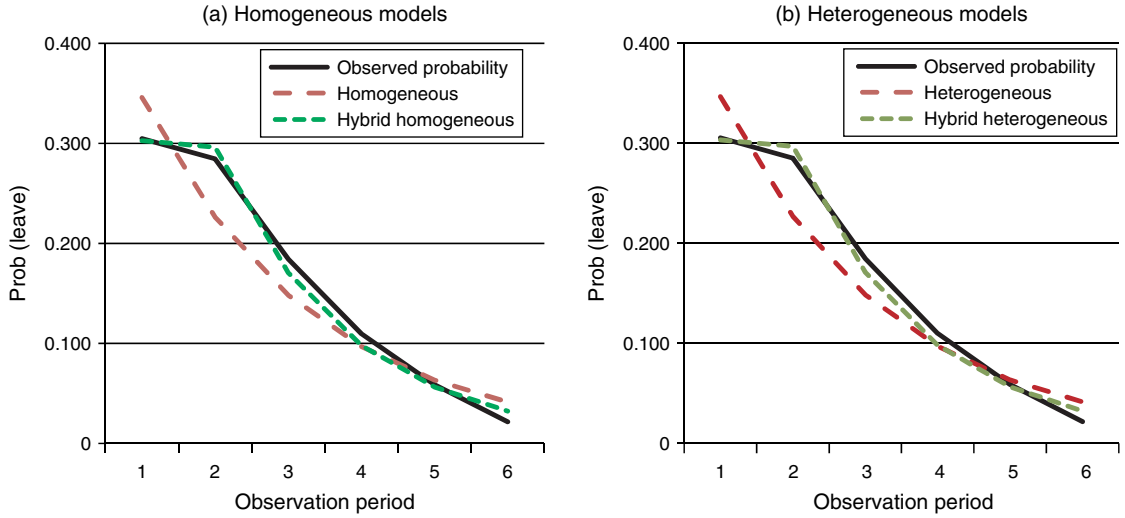
In November–December 2009 Suruga recruited customers from the Interface Asia panel. Screening and incentives were similar to those in the calibration study—the initial response rate was 22.1%. After screening on interest in card loans and age requirements, 10,182 out of 13,696 potential respondents were declared not eligible for the study. The remaining 3,514 were directed to the card-loan site, of which 1,997 explored the website for at least 2(1/2) minutes and 10 clicks (56.9%). Of these customers, 1,395 completed pre- and postvisit questionnaires providing valid data (70.1%). This is a net completion rate of 39.7% and an overall completion/response rate of 8.8%. The market research provider, Applied Marketing Science, Inc., reports that such overall rates are typical of complex Web-based studies. Of the 1,395 completions, 1,061 respondents experienced website morphing (a when-to-morph algorithm, test condition), and 334 experienced a static website (control condition). This was not a sufficient sample for either the HULB algorithm or the WTM algorithm to converge. However, it was sufficient to establish that it was feasible to decide when to morph within a customer's visit. On measures such as ease of use, providing relevant information, helpful to decisions, trustworthy, and recommendation, the morphing website did as well as the static website (no significant differences on 15 evaluative scales). Had the sample size been larger, perhaps the morphing website would actually have done better.

6.6. Analysis of the Random Time to Exit

We use the Suruga experience to examine which random-exit model would have fit the data best. For each respondent we observe the observation period in which the respondent left the website. These data are plotted in Figure 5 as a solid black line. The dashed red line plots the performance of a homogeneous model ($\psi_{tn} = \psi$ for all t, n, where ψ is determined by maximum-likelihood estimation). The homogeneous model explains 79.6% of the uncertainty ($U^2 = 0.796$), but visual inspection suggests a kink for the first observation period. To model the kink, we fit a hybrid homogeneous model ($\psi_{tn} = \psi_1$ if $t = 1$, but $\psi_{tn} = \psi$ otherwise) as shown by the green fine-dashed line in Figure 5(a). The fit is significantly better ($\chi_1^2 = 58.0$, $p < 0.01$) and explains almost all of the uncertainty ($U^2 = 0.968$).

We next examine the heterogeneous model described in §3.3.1. Using the functional invariance property of maximum-likelihood estimation, we reparameterized the model to estimate $E_n[\psi] = \alpha_\psi/(\alpha_\psi + \beta_\psi)$ and $Z_\psi \equiv (\alpha_\psi + \beta_\psi)$. The maximum-likelihood estimate of

Figure 5 **Model Comparison—Random Exit (ψ_{tn})**



$E_n[\hat{\psi}]$ is 0.346, however, the maximum-likelihood estimate of Z_ψ appears to diverge. The likelihood was still increasing at $\hat{Z}_\psi > 50{,}000$. The resulting heterogeneous model is indistinguishable from a homogeneous model ($\chi_1^2 \cong 0, p > 0.90$). We also fit a hybrid heterogeneous model with $\hat{\psi}_1 = 0.3027$, $E_n[\hat{\psi}] = 0.425$, and $\hat{Z}_\psi > 50{,}000$ for $t > 1$. The hybrid heterogeneous model was significantly different than the heterogeneous model ($\chi_1^2 = 58.0$, $p < 0.01$) and indistinguishable from the hybrid homogeneous model ($\chi_1^2 \cong 0$). The heterogeneous models are plotted in Figure 5(b).

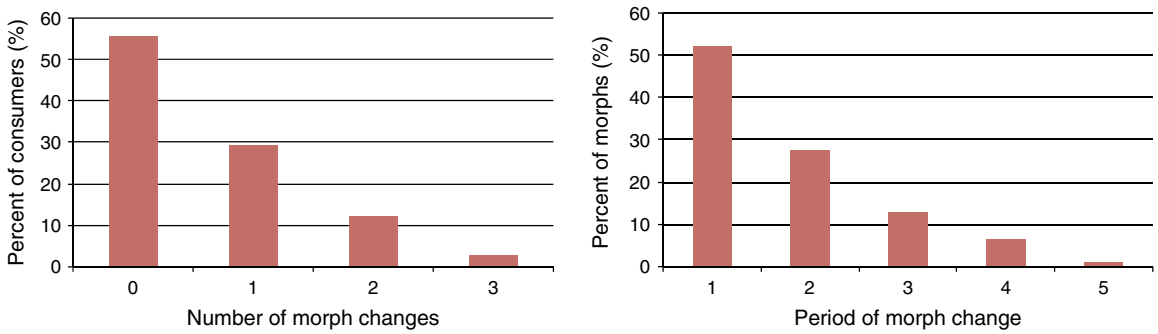
In the Suruga data, a hybrid homogeneous model of random exit appears to fit the data best. A fully time-varying model ($\psi_{tn} = \psi_t$) is significantly better than a hybrid homogeneous model ($\chi_4^2 = 11.0$, $p = 0.03$), but the slight increase in U^2 on a sample of 1,395 respondents may not be worth the sacrifice in parsimony. We estimated these models a posteriori to demonstrate that future applications could readily fit descriptive models using only the calibration data.

6.7. Descriptive Statistics for the Time to a Morph Change

If initial morphs were assigned randomly, then we would expect the initial morph to be optimal for roughly 25% of the respondents. However, both the HULB and WTM algorithms assign the initial morph based on the expected Gittins indices using the prior probabilities that customer n belongs to each segment. If an algorithm works well, we expect the initial morphs to be optimal for more than 25% of the respondents, especially when one segment is larger than the others. (For Suruga, the largest segment was approximately 43.2% of the respondents.) Furthermore, to avoid switching costs, the WTM algorithm may keep the initial morphs if the expected improvement in the expected Gittins index is small.

Performance was consistent with these expectations. The implemented WTM algorithm stayed with the initial morph roughly half of the time. Of those respondents who experienced morph changes, most experienced a single change; very few experienced three or more changes (Figure 6). Figure 6 also suggests that, empirically, the WTM algorithm was able to identify the morph changes rapidly. The majority of morph changes were made after the first period, with decreasing numbers in subsequent periods. This is consistent with a website designed explicitly to identify customer segments rapidly.

Figure 6 **Suruga Bank Experiment: Descriptive Statistics**



6.8. Attempted Comparison Between a WTM Algorithm and a Static Website

To evaluate performance, we collected measures that were not maximized by the algorithm: consideration, preference, and purchase likelihood. Consideration was a binary consider-or-not measure, preference was a 100-point constant sum scale, and purchase likelihood was measured on an 11-point scale as in Juster (1966). All measures reflect that a customer will not prefer or purchase from Suruga Bank if they do not first consider Suruga Bank. We also attempted to correct for premeasures, but because Suruga Bank was not known for card loans at the time of the application, only 70 respondents considered Suruga Bank in the premeasures. Unfortunately, because the sample size was small relative to the synthetic-data experiments and Urban et al. (2014), none of the measures were significantly different between the test and control cells. The results are available in the supplemental material.

6.9. Counterfactual Policy Simulations with Suruga Morph × Segment Purchase Probabilities

Although the Suruga Bank experiment ran for only 1,395 customers, we can run a counterfactual policy simulation in which there are 40,000 customers in each of three experimental cells—WTM, HULB, and random assignment. We use the same code that reproduced the HULB simulations (Table 1, §5.1) except that we use the $\hat{p}_{rm}$ from the Suruga Bank implementation. Table 4 reports the results. HULB is substantially better than random assignment, and WTM is substantially better than HULB. In this case, the improvement due to WTM (63%) is almost 5× that of HULB. WTM does well because, for the counterfactual policy simulation, it is better to morph much earlier than the HULB default. The $\hat{p}_{rm}$ differ among morphs more for the Suruga Bank website than they do for the BT Group website. Although a devil's advocate might argue that a revised earlier-morphing HULB might perform better than the tested HULB, such an insight was made possible by the dynamic program in the WTM algorithm and would be hard to know a priori without the analyses in this paper.

7. Summary and Challenges

The HULB website-morphing algorithm has proven to increase outcomes with both synthetic data (BT Group application) and in the field (Urban et al. 2014), but it ignores switching costs, the impact of premorph clicks, and random exit. The WTM algorithm generalizes the HULB algorithm to allow more realistic assumptions. The WTM algorithm runs in real time between clicks on a website. Synthetic data analyses of two data sets suggest that the WTM algorithm can improve outcomes substantially if the tuning parameters are set properly. The algorithm appears to be reasonably robust to misspecification of the tuning parameters. Furthermore, the Suruga Bank application suggests that it is feasible on a real-world website to determine when to morph in real time.

There remain many challenges. (1) Large-sample empirical applications might (a) parse the incremental value of the WTM algorithm relative to the special case of HULB and (b) calibrate the values of the switching-cost and impact-weight tuning parameters. (2) Improved theory might (a) establish optimality properties for the linkage between the morph-to-segment and when-to-morph optimizations and (b) improve "fractional-observation" updating of the posterior distributions of the p_{rmn}'s. (3) Improved algorithms might (a) relax assumptions of multiplicative switching costs, additive impact weights, and independent website exit; (b) model network externalities such as the ability of customers to share knowledge of a website's look and feel; (c) allow for switching costs that depend on the multiple morphs ($\gamma_{mm'}$); and (d) allow exit probabilities that depend on segments, final morphs, or morph history (ψ_{trm}, $\psi_{trmm'}$, etc.).

Table 4 Counterfactual Reanalysis of Suruga Bank Experiment If Run to Convergence Using the When-to-Morph Algorithm and Observed Morph × Segment Probabilities

	Reward at convergence[a]	Percent improvement in reward[b] (%)	Net present value (NPV)[c]	Percent improvement in NPV[d] (%)
Baseline for no morphing	0.0632	0	0.0619	0
Website morphing				
Website morphing (HULB)	0.0688	13	0.0667	12
Website morphing with improved algorithm	0.0912	66	0.0882	63
Perfect information ($n \to \infty$)[e]	0.1055	100	0.1034	100

Notes. Switching costs are $\gamma = 0.95$. Impact weights are $\vec{w} = (0.25, 0.25, 0.25, 0.25)$. $\hat{\Omega}$ is as in Table 1.

[a]Rewards for the last 400 customers out of 40,000 customers.

[b](Reward due to algorithm − reward at no morphing)/(reward due to perfect information − reward at no morphing).

[c]Net present value of the rewards divided by the number of periods.

[d](NPV due to algorithm − NPV at no morphing)/(NPV due to perfect information − NPV at no morphing).

[e]This is the upper bound. Applications do not have perfect information on either customer segments or purchase probabilities.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/mnsc.2014.1961.

Acknowledgments

The authors gratefully acknowledge the contributions of their industrial collaborators, research assistants, and faculty colleagues: Shirley S. Fung, Steven Gaskin, Robert Klein, Clarence Lee, Qiuyuan Jimmy Li, Erin Mac-Donald, Eleojo Ocholi, Chris Perciballi, Takuya Sugiyama, Yoshio Tokoro, and Kevin Wang. They appreciate the comments of participants at the Erasmus Research Camp, New England Marketing Research Conference, Marketing Science Conference, and seminars at the Australian National University and the Wharton School. This research was supported by the MIT Sloan School of Management, the Center for Digital Business at MIT (ebusiness.mit.edu), Suruga Bank, Applied Marketing Science, Inc., and the Erasmus Center for Marketing and Innovations.

Appendix A. Notation Used in This Paper and in Appendix B

- a The amount by which we discount a sale to customer $n+1$ relative to customer n
- $B(\alpha, \beta)$ Beta distribution with parameters α and β
- c_{tnj} Indicator variable to indicate whether customer n chooses click alternative j at t
- c_{tn} Customer n's clicks in the tth observation period; n may be suppressed
- $\vec{c}_{tn}$ Customer n's clicks up to and including the tth observation period, $\{c_{1n}, c_{2n}, \ldots, c_{tn}\}$.
- EG_{mn} Expected Gittins index for mth morph for customer n
- EIR Expected immediate reward (used in the when-to-morph Bellman equation)
- $f(\cdot \mid \cdot)$ Posterior distribution; specific distribution when arguments are specific; sometimes subscripted by n
- G_{rmn} Gittins index for rth segment and mth morph for customer n; n may be suppressed
- HULB Hauser, Urban, Liberali, and Braun (2009)
- j Indexes click alternatives (links)
- J_{tn} The number of click alternatives faced by customer n on the tth click
- m Indexes morphs
- m_{tn} Morph that customer n sees in observation period t
- m^*_{tn} Optimal morph for customer n in observation period t
- M Number of morphs available from which to choose
- n Indexes customers
- n_{ob} Number of customers in an observation period; used in Appendix B
- $n_{ob,\,rm}$ Number of consumers in an observation period who were in segment r and saw morph m
- p_{rmn} Probability that customer n in segment, r, purchases when shown morph m
- $\Pr_0(r_n = r)$ Prior probability that customer n belongs to segment r
- $q_{rn}(\vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n})$ Probability that customer n is in segment r
- q_{rn} Shorthand for $q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})$; used in Appendix B
- $q_{rn}(s)$ Shorthand for $q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n} \mid r_{\text{true}} = s)$; used in Appendix B
- r Indexes customer segments
- R Number of customer segments
- $r_{\text{true},\,n}$ Customer n's true segment (used derive the when-to-morph dynamic program)
- s Used in summations
- S Observation period (used in random exit equation)
- S_s Set of customers in the last n_{ob} customers whose true segment is s; used in Appendix B.
- t Indexes observation periods
- T_n Number of observation periods for customer n
- $\tilde{u}_{tnj}$ Utility customer n obtains from clicking on the jth click alternative on the tth click
- $V_t(m^*_{tn}, m_{t-1,\,n}, \vec{c}_{t-1,\,n}, \hat{\Omega}, \vec{X}_{t-1,\,n}, N_{t-1,\,n})$ Continuation value function for when-to-morph decision
- $V_{\text{Gittins}}(\alpha_{rmn}, \beta_{rmn}, a)$ Continuation value function when calculating the Gittins index
- w_t Impact weight for the tth observation period
- $\vec{w}$ Vector of the w_t's; $\hat{\vec{w}}$ if estimated from the calibration data
- WTM When-to-morph algorithm; proposed generalization of the HULB algorithm
- $\vec{x}_{tnj}$ Vector of characteristics of the jth click alternative up to the tth click of customer n
- $\vec{X}_{tn}$ Set of all $\vec{x}_{tnj}$'s for all clicks up to an including click t, $\{\vec{x}_{1nj}, \vec{x}_{2nj}, \ldots, \vec{x}_{tnj}\}$ $\forall\, j = 1$ to J_{tn}
- Z_ψ Equal to $\alpha_\psi + \beta_\psi$
- $\hat{Z}_\psi$ Maximum-likelihood estimate of Z_ψ
- α_{rmn} Parameter of the beta distribution used to model uncertainty over p_{rmn}
- β_{rmn} Parameter of the beta distribution used to model uncertainty over p_{rmn}
- α_ψ Parameter of a beta distribution used to model heterogeneity in exit probabilities
- β_ψ Parameter of a beta distribution used to model heterogeneity in exit probabilities
- γ Discount for switching from one morph to another; after a switch, the probability of a purchase is reduced by a factor of γ ($\hat{\gamma}$ if estimated from the calibration data)
- δ_n Indicator variable to indicate whether customer n makes a purchase; used in the when-to-morph algorithm when multiple morphs (might) affect outcomes
- δ_{nm} Indicator variable to indicate whether customer n makes a purchase given that customer n saw morph m after a switch; used in HULB
- δ_{rmn} Indicator variable to indicate whether customer n makes a purchase given that customer n is in known segment r and saw morph m; used in Appendix B.
- $\Delta_{m'_{tn}tn}$ Indicator variable to indicate if we change to morph m'_{tn} in period t for customer n
- $E_n[\hat{\psi}]$ Maximum likelihood estimate of $\alpha_\psi/(\alpha_\psi + \beta_\psi)$
- $\tilde{\epsilon}_{tnj}$ Extreme value error; used to model customer n's preference for click alternatives

η_{mnt} Indicator variable to indicate if customer n saw morph m in observation period t
Σ Used to indicate summation, not a variable
τ Used as a variable for summation and for ranges of t
τ_o In HULB, the number of clicks observed prior to a morph
ψ_{tn} Probability that customer n does not continue in the tth observation given the customer was on the website in the $t+1$st observation period
$\Psi_n(S \mid t-1)$ Probability customer n is still on the website in observation period S given the customer was on the website in observation period $t-1$
$\bar{\Psi}(S \mid t-1)$ Expectation of $\Psi_n(S \mid t-1)$ over customers
$\vec{\omega}_r$ Vector of preference weights for the $\vec{x}_{tnj}$ for the rth segment
Ω Matrix of the $\vec{\omega}_r$
$\hat{\Omega}$ Estimate of Ω (either maximum likelihood or Bayesian mean posterior)
ζ Used as a variable for summation in Appendix B to sum over customers beyond n

Appendix B. Full-Probability Model and Updating Formulae

B.1. Review of Updating for Gittins Indices

Notation follows the text. If we knew the customer's segment, r, and the customer saw the same morph, m, for the entire visit, then the relevant posterior probability after customer n would be p_{rmn}, the probability that customer n in segment r who saw morph m would make a purchase. We assume that the prior distribution for p_{rmn} is beta as given by

$$f_n(p_{rmn} \mid \alpha_{rmn}, \beta_{rmn}) = B(\alpha_{rmn}, \beta_{rmn}) p_{rmn}^{\alpha_{rmn}-1}(1-p_{rmn})^{\beta_{rmn}-1}, \quad \text{(B1)}$$

where α_{rmn} and β_{rmn} are parameters of the beta distribution, and $B(\alpha_{rmn}, \beta_{rmn})$ is the beta function.

We now observe whether or not customer n makes a purchase. Let $\delta_{rmn} = 1$ if customer n in *known* segment r made a purchase and $\delta_{rmn} = 0$ otherwise. The data likelihood is then given by a binomial distribution:

$$f_n(\delta_{rmn} \mid p_{rmn}) = p_{rmn}^{\delta_{rmn}}(1-p_{rmn})^{1-\delta_{rmn}}. \quad \text{(B2)}$$

Combining Equations (B1) and (B2), the posterior distribution for $p_{rm,\,n+1}$ is proportional to

$$f_{n-1}(p_{rm,\,n+1} \mid \alpha_{rmn}, \beta_{rmn}, \delta_{rmn}) \sim p_{rm,\,n+1}^{\alpha_{rmn}+\delta_{rmn}-1}(1-p_{rm,\,n+1})^{\beta_{rmn}+(1-\delta_{rmn})-1}. \quad \text{(B3)}$$

Equation (B3) is again a beta distribution. Thus we have the updating equation as given by

$$\begin{aligned} \alpha_{rm,\,n+1} &= \alpha_{rmn} + \delta_{rmn}, \\ \beta_{rm,\,n+1} &= \beta_{rmn} + (1-\delta_{rmn}). \end{aligned} \quad \text{(B4)}$$

We easily extend the formal updating to the case where we observe n_{ob} additional customers before we update (read n_{ob} as "n observed"). In this case we have a binomial data likelihood and we allow a more general definition of $\delta_{rmn_{ob}}$ as the number of purchases made by the n_{ob} customers that follow the nth customer. Following similar steps, the posterior distribution is also a beta distribution with parameters given by

$$\begin{aligned} \alpha_{rm,\,n+n_{ob}} &= \alpha_{rmn} + \delta_{rmn_{ob}}, \\ \beta_{rm,\,n+n_{ob}} &= \beta_{rmn} + (n_{ob} - \delta_{rmn_{ob}}). \end{aligned} \quad \text{(B5)}$$

Because updating for the Gittins index is naturally conjugate, computations are simple and rapid. It is feasible to update after every customer or observation period. Furthermore, by comparing Equations (B4) and (B5) it is obvious that we obtain the same posterior distribution at the $(n+n_{ob})$th customer whether we update sequentially using Equation (B5) or update all n_{ob} customers at one time.

B.2. Review of Updating for the HULB Hidden Markov Model

HULB extend Gittins updating to a situation where we do not know the customer's segment, but rather have estimates, $q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})$, that customer n belongs to segment r for each possible segment. HULB continue to assume that the customer saw a single morph, m, for a sufficient fraction of the visit that we need only consider the effect of morph m on customer n's purchase. With these assumptions, the probability, p_{mn}, that customer n who saw morph m makes a purchase is given by

$$p_{mn} = \sum_{r=1}^{R} q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n}) p_{rmn}. \quad \text{(B6)}$$

HULB assume a beta prior distribution as in Equation (B1). For indexability, HULB assume that each r, m-arm of the bandit process is independent. Independent arms imply that the prior distributions are independent over r and m. The joint prior distribution generalizes Equation (B1) to become Equation (B7) prior to observing δ_{mn}. For simplicity of notation, we let $\vec{\alpha}_{mn}$ and $\vec{\beta}_{mn}$ be the vectors of parameters (over r):

$$f_n(p_{rmn} \mid \vec{\alpha}_{mn}, \vec{\beta}_{mn}) = \prod_{r=1}^{R} B(\alpha_{rmn}, \beta_{rmn}) p_{rmn}^{\alpha_{rmn}-1}(1-p_{rmn})^{\beta_{rmn}-1}. \quad \text{(B7)}$$

The data likelihood is binomial and generalizes Equation (B2):

$$\begin{aligned} f_n(\delta_{rmn} \mid p_{rmn}) &= p_{mn}^{\delta_{mn}}(1-p_{mn})^{1-\delta_{mn}} \\ &= \left(\sum_{r=1}^{R} q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n}) p_{rmn}\right)^{\delta_{mn}} \\ &\quad \cdot \left(1 - \sum_{r=1}^{R} q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n}) p_{rmn}\right)^{1-\delta_{mn}}. \end{aligned} \quad \text{(B8)}$$

We extend the analysis to allow updating after n_{ob} customers. Unlike in the single-armed case, the analysis does not simplify. Specifically, each customer has a different set of $q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})$ values based on that customer's unique clickstream. Thus, the data likelihood is now given by Equation (B9), where ζ is notation for use in the product so that it does not conflict with n or n_{ob}:

$$\begin{aligned} &f_{n+n_{ob}}(\delta_{rmn_{ob}} \mid p_{rm,\,n_{ob}}) \\ &= \prod_{\zeta=n+1}^{n+n_{ob}} \left\{ \left(\sum_{r=1}^{R} q_{r\zeta T_\zeta}(\vec{c}_{T_\zeta \zeta}, \hat{\Omega}, \vec{X}_{T_\zeta \zeta}) p_{rm,\,n+n_{ob}}\right)^{\delta_{m\zeta}} \right. \\ &\quad \left. \cdot \left(1 - \sum_{r=1}^{R} q_{r\zeta T_\zeta}(\vec{c}_{T_\zeta \zeta}, \hat{\Omega}, \vec{X}_{T_\zeta \zeta}) p_{rm,\,n+n_{ob}}\right)^{1-\delta_{m\zeta}} \right\}. \end{aligned} \quad \text{(B9)}$$

Equation (B9) uses the substitution in Equation (B6) for $p_{rm,\,n+n_{ob}}$ for all ζ in the product. This implies that the posterior distribution after observing n_{ob} additional customers is given by Equation (B10):

$$f_{n+n_{ob}}(p_{rm,\,n+n_{ob}} \mid \vec{\alpha}_{mn}, \vec{\beta}_{mn}, \delta_{rmn_{ob}}) \sim \prod_{\zeta=n+1}^{n+n_{ob}} \left\{ \left(\sum_{r=1}^{R} q_{r\zeta T_\zeta}(\vec{c}_{T_\zeta\zeta}, \hat{\Omega}, \vec{X}_{T_\zeta\zeta}) p_{rm,\,n+n_{ob}} \right)^{\delta_{m\zeta}} \cdot \left(1 - \sum_{r=1}^{R} q_{r\zeta T_\zeta}(\vec{c}_{T_\zeta\zeta}, \hat{\Omega}, \vec{X}_{T_\zeta\zeta}) p_{rm,\,n+n_{ob}} \right)^{1-\delta_{m\zeta}} \right\} \cdot \prod_{r=1}^{R} p_{rm,\,n+n_{ob}}^{\alpha_{rmn}-1} (1 - p_{rm,\,n+n_{ob}})^{\beta_{rm,\,n}-1}. \tag{B10}$$

B.2.1. The Infeasibility of Using This Posterior Distribution for Website Morphing. If we knew every customer's segment, the $q_{r\zeta T_\zeta}(\vec{c}_{T_\zeta\zeta}, \hat{\Omega}, \vec{X}_{T_\zeta\zeta})$ would be either 0 or 1, and Equation (B10) would become naturally conjugate. When the $q_{r\zeta T_\zeta}(\vec{c}_{T_\zeta\zeta}, \hat{\Omega}, \vec{X}_{T_\zeta\zeta})$ are 0 or 1, we obtain the same updating formulae as in Equation (B5).

When the segments are not known with certainty, the $q_{r\zeta T_\zeta}(\vec{c}_{T_\zeta\zeta}, \hat{\Omega}, \vec{X}_{T_\zeta\zeta})$ are not 0 or 1, and we have the full summation in Equation (B10). We no longer have naturally conjugate updating; hence, $f_{n+n_{ob}}(n_{ob} \mid \vec{\alpha}_{mn}, \vec{\beta}_{mn}, \delta_{rmn_{ob}})$ is not a beta distribution, even when $n_{ob} = 1$. If we were only interested in the posterior distribution of $p_{rm,\,n+n_{ob}}$, it might be feasible to discretize $p_{rm,\,n+n_{ob}}$ and use Equation (B10) to compute numerically the marginal posterior distribution for $p_{rm,\,n+n_{ob}}$, perhaps after integrating out the $p_{rm',\,n+n_{ob}}$ values for $m' \neq m$. Computations will not be rapid between customers on a high-traffic website, but it might be feasible to compute the posterior distributions at the end of each day and use the updated distributions for the following day.

However, even if we chose to compute the posterior distribution offline, we must use the posterior distributions to compute the indices (via a dynamic program) to assign morphs. HULB exploit the property that the Gittins index for beta-binomial updating is particularly simple and can be easily tabled for α_{rmn} and β_{rmn}. They table the indices based on updated α_{rmn} and β_{rmn} and look them up as needed. It is not feasible to resolve Gittins' (1979) dynamic program after each customer. Morphing when customer segments follow a hidden Markov model is indexable, as proven by Krishnamurthy and Mickova (1999), and indexability does not depend on the fact that the posterior is an analytic distribution. Unfortunately, we know of no way to compute the index sufficiently rapidly when the posterior distribution is defined numerically.

B.2.2. Practical Solution: Fractional Observations. To obtain a feasible solution, HULB make a heroic independence assumption and use the concept of fractional observations. We first note that the binomial distribution, a probability mass function, can also be interpreted as a probability density function for fractional observations. We then interpret the outcome variable, δ_{mn}, as a combination of fractional outcomes for each potential customer segment, r, and treat the fractional outcomes as if they were independent over r. In particular, the fractional outcomes become $\delta_{rmn} = q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})\delta_{mn}\ \forall r$. Using this interpretation and the shorthand notation $q_{rn} = q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})$, the data likelihood becomes[2]

$$f_n(q_{rn}\delta_{mn} \mid p_{rmn}) = p_{rmn}^{q_{rn}\delta_{mn}} (1 - p_{rmn})^{q_{rn}(1-\delta_{mn})}. \tag{B11}$$

We continue to use Equation (B1) as the prior distribution, which gives Equation (B12) as the posterior distribution:

$$f_{n+1}(p_{rm,\,n+1} \mid \alpha_{rmn}, \beta_{rmn}, q_{rn}\delta_{rmn}) \sim p_{rm,\,n+1}^{\alpha_{rmn}+q_{r,\,n}\delta_{mn}-1} (1 - p_{rm,\,n+1})^{\beta_{rmn}+q_{rn}(1-\delta_{mn})-1}. \tag{B12}$$

With fractional observations, the posterior distribution is naturally conjugate, and the updating equations become (HULB, Equation (2), p. 210), using full notation,

$$\begin{aligned} \alpha_{rm,\,n+1} &= \alpha_{rmn} + q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})\delta_{mn}, \\ \beta_{rm,\,n+1} &= \beta_{rmn} + q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})(1 - \delta_{mn}). \end{aligned} \tag{B13}$$

B.3. Extending the HULB Updating Formula to Switching Costs, Impact Weights, and Multiple Morphs

Once we accept the concept of fractional observations as a practical means to retain the structure of expected Gittins index updating, we can extend Equation (B13) to account for switching costs, impact weights, and multiple morphs. No new derivations need to be introduced. We simply substitute the new fractional observations into Equations (B12)–(B15). The new concepts are that the switching costs, impact weights, and multiple morphs may affect customer n's purchase probabilities.

First, customer n does not see morph m for the entire website visit; customer n sees morph m for those periods in which the morphing algorithm selects morph m. The effect on purchasing is proportional to the impact weights for the periods in which morph m was shown. As defined in the text, $\eta_{mnt} = 1$ if morph m is shown to customer n in observation period t, and w_t is the impact weight for observation period t. The probability is also affected by whether or not a switch in morphs occurs in the period as summarized by $\Delta_{m_{tn}tn}$. The new fractional observation becomes $q_r(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})(\sum_{t=1}^{T} \gamma^{\Delta_{m_{tn}tn}} \eta_{mnt} w_t)$.

With the revised definition of a fractional observation, the revised updating formula becomes Equation (B14), which is Equation (9) in the text. The data likelihood and the posterior distributions generalize Equations (B11) and (B12) as expected:

$$\begin{aligned} \alpha_{rm,\,n+1} &= \alpha_{rmn} + q_r(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})\gamma^{N_{T_n}} \left(\sum_{t=1}^{T_n} \eta_{mnt} w_t \right) \delta_n, \\ \beta_{rm,\,n+1} &= \beta_{rmn} + q_r(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})\gamma^{N_{T_n}} \left(\sum_{t=1}^{T_n} \eta_{mnt} w_t \right) \cdot (1 - \delta_n). \end{aligned} \tag{B14}$$

The updating formulae in Equation (B14) are practical because they maintain a naturally conjugate beta posterior distribution, because it can be computed quickly, and because, using arguments similar to those in §B.4, it should converge to the true values, $p_{rm,\,\text{true}}$, when the assumed model of behavior is correct.

[2] We examine in §2.4 why the updating formulae are still likely to converge, and we present evidence of convergence from the synthetic data experiments.

B.4. Motivation of Why Fractional Observations Are a Good Approximation

The trick of fractional observations requires that we assume a level of independence that we know is violated. Despite this assumption, both the HULB algorithm and the WTM algorithm improve outcomes substantially with synthetic data—the HULB algorithm relative to a static website and the WTM algorithm relative to the HULB algorithm. In this section we motivate why that is likely that fractional observations are a good approximation. Our arguments are not a formal proof, and they do not rule out the algorithms getting trapped in local minima; rather they suggest why we expect the fractional-observation updating formulae to converge to $p_{rm,\text{true}}$.

For simplicity of notation, consider the HULB assumptions of $\gamma = 1$ and $\vec{w} = (0, 0, 0, 1)$. We are interested situations where the data overwhelm the priors; that is, assume we observe $n_{ob,m}$ customers where $n_{ob,m}$ is sufficiently large: $n_{ob,m} \gg \sum_{r=1}^{R}(\alpha_{rm1} + \beta_{rm1})$. Let S_{rm} be the set of customers whose true segment is r and who saw morph m. We continue to use the shorthand notation, $q_{rn} = q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n})$, and define $q_{rn}(s) = q_{rnT_n}(\vec{c}_{T_n n}, \hat{\Omega}, \vec{X}_{T_n n} \mid r_{\text{true}} = s)$. Under the conditions that the data overwhelm the priors, the HULB updating formulae (Equation (B13)) become (for all r and m)

$$\begin{aligned}\alpha_{rmn_{ob}} &= \sum_{n=1}^{n_{ob,m}} q_{rn}\delta_{mn} = \sum_{s=1}^{R}\sum_{n\in S_{sm}} q_{rn}(s)\delta_{mn},\\ \beta_{rmn_{ob}} &= \sum_{n=1}^{n_{ob,m}} q_{rn}(1-\delta_{mn}) = \sum_{s=1}^{R}\sum_{n\in S_{sm}} q_{rn}(s)(1-\delta_{mn}).\end{aligned} \tag{B15}$$

We now take expected values:

$$\begin{aligned}E_n[\alpha_{rmn_{ob}}] &= \sum_{s=1}^{R}\sum_{n\in S_{sm}} q_{rn}(s)p_{sm,\text{true}},\\ E_n[\beta_{rmn_{ob}}] &= \sum_{s=1}^{R}\sum_{n\in S_{sm}} q_{rn}(s)(1-p_{sm,\text{true}}).\end{aligned} \tag{B16}$$

Under the assumption that customers are homogeneous within true segments such that $q_{rn}(s) = q_r(s)$, we recognize that

$$\begin{aligned}E_n[\alpha_{rmn_{ob,m}}] &= \sum_{s=1}^{R} q_r(s)n_{ob,sm}p_{sm,\text{true}} = q_r(r)n_{ob,rm}p_{rm,\text{true}}\\ &\quad + \sum_{s\neq r} q_r(s)n_{ob,sm}p_{sm,\text{true}},\\ E_n[\beta_{rmn_{ob,m}}] &= \sum_{s=1}^{R} q_r(s)n_{ob,sm}(1-p_{sm,\text{true}})\\ &= q_r(r)n_{ob,rm}(1-p_{rm,\text{true}})\\ &\quad + \sum_{s\neq r} q_r(s)n_{ob,sm}(1-p_{sm,\text{true}}).\end{aligned} \tag{B17}$$

If the first term dominates (the term containing $q_r(r)$), the beta posterior will converge to a point mass at $p_{rm,\text{true}}$ as $n_{ob,m}$ gets large. It will converge to a point mass because both $E_n[\alpha_{rmn_{ob}}]$ and $E_n[\beta_{rmn_{ob}}]$ are the order of $n_{ob,m}$. Furthermore, $p_{rmn_{ob,n}} = E[\alpha_{rmn_{ob,m}}]/\{E[\alpha_{rmn_{ob,m}}] + E[\beta_{rmn_{ob,m}}]\}$, which approach $p_{rm,\text{true}}$ when the first term dominates (because $q_r(r)n_{ob,rm}$ drops out).

There are two forces that make it likely that, as $n_{ob,m}$ gets large, the first term is much larger than the remaining terms for the morphs that are chosen for segment r. First, if $\hat{\Omega}$ is well estimated and the logit model does a good job in predicting customer segments, then $q_r(r) \gg q_r(s)\ \forall s \neq r$. Second, as the algorithm gets better at identifying the best morph for segment r (and if there is sufficient variation in morphs that different morphs are best for different segments), then $n_{ob,rm_r^*} \gg n_{ob,sm_r^*}$ for $s \neq r$. (m_r^* is the best morph for segment r.) If the algorithm is indeed close to optimal, m_r^* will be chosen for segment r. Convergence will be best for optimal segment x morph matches, and especially for those morphs that are shown many times. (Because it is likely that $p_{sm_r^*} \leq p_{rm_r^*}$, we can also argue that convergence will be from below. Details are available in the supplemental material.)

It is an empirical question whether these two forces are sufficiently strong so that fractional updating converges. Fortunately, synthetic-data experiments and empirical experience suggest that the HULB algorithm and the generalized algorithm, which also relies on fractional-observation updating, lead to improved outcomes and the ability to match morphs to segments. For the synthetic-data experiments in §5, we can compare the true purchase probability values, $p_{rm,\text{true}}$ values, to the posterior means. With $\gamma = 1$ and $\vec{w} = (0, 0, 0, 1)$, we test the HULB updating formulae in Equation (B13). For this case, the mean absolute percent error is approximately 0.028, which is 8.7% of the average true probability. As the number of customers gets large, the algorithm tends toward identifying the optimal morph for each segment; hence we gain the most information about optimal segment × morph matches (r and m_r^*). When we look at optimal segment × morph matches, the error reduces to 5.8% of the average true probability. This reduces to 0.5% when we require a morph to have been shown to customers at least 2,000 times.

We examine Equation (B14) with $\gamma = 0.95$ and $\vec{w} = (0.25, 0.25, 0.25, 0.25)$. For this case, the mean absolute percentage error is approximately 0.0318, which is 10.0% of the average true probability. When we look at the optimal morph for each segment, the error reduces to 7.3% of the average true probability. This reduces to 3.4% when we require a morph to have been shown at least the equivalent of 2,000 times (i.e., $2{,}000T_n$ total periods). The mean absolute error decreases further when morphs are shown even more often. We see similar patterns for $\gamma = 0.80$ and $\vec{w} = (0.25, 0.25, 0.25, 0.25)$.

Appendix C. Supplemental Material Available at http://dx.doi.org/10.1287/mnsc.2014.1961

C.1. Synthetic Data and Counterfactual Code and Pseudocode (Tables 1, 2, 4; §5, §6.9)

- Pseudocode and *R* code with which to run the 40,000 synthetic-customer simulations using specified values for $\hat{p}_{rm}$, γ, and $\vec{w}$
- Spreadsheet with the values of $\hat{p}_{rm}$ for the BT Group simulations and the Suruga Bank counterfactual policy simulation

C.2. WTM Code for Suruga Bank Application (§6)

- *R* code (that reproduces original PHP code) that was called by the Suruga Bank website to determine the morph-to-segment match and when to morph

C.3. Data from Suruga Bank Experiment (Figure 6, §6)

- Table comparing insignificant measures of consideration, preference, and purchase likelihood
- Excel spreadsheet for comparison table
- Excel spreadsheet that produced the descriptive statistics (Figure 6)
- Excel spreadsheet for site-evaluation measures

C.4. Japanese with English Translations of Suruga Bank Questionnaires (§6)

- Japanese with English translation of the questionnaire used in the calibration study
- Japanese with English translation of the questionnaire used to obtain evaluative measures for the Suruga Bank experiment

C.5. Demonstration That Fractional Updating Will Likely Converge from Below (Appendix B)

- Short derivation to illustrate why fractional updating will likely converge from below

C.6. User's Guide to Implementation of Morphing, Including When to Morph

- A user's guide on how to implement morphing including both morph-to-segment assignment and the decision on when to morph; user's guide relies on citations to this paper, HULB, and Urban et al. (2014) to provide a practical roadmap to morphing.

References

Asawa M, Teneketzis D (1996) Multi-armed bandits with switching probabilities. *IEEE Trans. Automatic Control* 41(3):328–348.

Balabanis G, Reynolds N, Simintiras A (2006) Bases of e-store loyalty: Perceived switching barriers and satisfaction. *J. Bus. Res.* 59(2):214–224.

Banks JS, Sundaram RK (1994) Switching costs and the Gittins index. *Econometrica* 62:687–694.

Bertsimas D, Niño-Mora J (1996) Conservation laws, extended polymatroids and multi-armed bandit problems. *Math. Oper. Res.* 21(2):257–306.

Dusonchet F, Hongler MO (2006) Priority index heuristic for multi-armed bandit problems with set-up costs and/or set-up time delays. *Internat. J. Comput. Integrated Manufacturing* 19(3):210–219.

Gittins JC (1979) Bandit processes and dynamic allocation indices. *J. Royal Statist. Soc., Ser. B (Methodol.)* 41(2):148–177 [plus commentary].

Gittins JC, Glazebrook K, Weber R (2011) *Multi-Armed Bandit Allocation Indices* (John Wiley & Sons, London).

Hauser JR, Toubia O (2005) The impact of utility balance and endogeneity in conjoint analysis. *Marketing Sci.* 24(3):498–507.

Hauser JR, Urban GL, Liberali G, Braun M (2009) Website morphing. *Marketing Sci.* 28(2):202–224.

Jersild AT (1927) Mental set and shift. *Arch. Psych.* 21(89):1–92.

Johnson EJ, Bellman S, Lohse GL (2003) Cognitive lock-in and the power law of practice. *J. Marketing* 67(2):62–75.

Jones MA, Mothersbaugh DL, Beatty SE (2000) Switching barriers and repurchase intentions in services. *J. Retailing* 76(2):259–274.

Jones MA, Mothersbaugh DL, Beatty SE (2002) Why customers stay: Measuring the underlying dimensions of services switching costs and managing their differential strategic outcomes. *J. Bus. Res.* 55:441–450.

Jun T (2004) A survey on the bandit problem with switching costs. *De Economist* 152(4):513–541.

Juster FT (1966) Consumer buying intentions and purchase probability: An experiment in survey design. *J. Amer. Statist. Assoc.* 61:658–696.

Krishnamurthy V, Mickova J (1999) Finite dimensional algorithms for the hidden Markov model multi-armed bandit problem. *IEEE Internat. Conf. Acoustics, Speech, and Signal Processing*, Vol. 5 (IEEE, New York), 2865–2868.

Little JDC (1966) A model of adaptive control of promotional spending. *Oper. Res.* 14(6):1075–1097.

Little JDC (1970) Managers and models: The concept of a decision calculus. *Management Sci.* 16(8):B466–B485.

Liu Q, Otter T, Allenby GM (2007) Investigating endogeneity bias in marketing. *Marketing Sci.* 26(5):642–650.

Meiran N (2000) Modeling cognitive control in task-switching. *Psych. Res.* 63:234–249.

Ny JL, Feron E (2006) Restless bandits with switching costs: Linear programming relaxations, performance bounds and limited lookahead policies. *Proc. 2006 Amer. Control Conf. of the Amer. Automatic Control Council, Troy, NY*, 1587–1592.

Spector A, Biederman I (1976) Mental set and shift revised. *Amer. J. Psych.* 89:669–679.

Tokoro Y (2008) Improving customer experience through advocacy and morphing: A Web application for Suruga Bank. Unpublished S.M. thesis, MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge.

Tsitsiklis JN (1994) A short proof of the Gittins index theorem. *Ann. Appl. Probab.* 4(1):194–199.

Urban GL, Liberali G, MacDonald E, Bordley R, Hauser JR (2014) Morphing banner advertisements. *Marketing Sci.* 33(1)27–46.

Weber R (1992) On the Gittins index for multi-armed bandits. *Ann. Appl. Probab.* 2(4):1024–1033.

Weiss AM, Anderson E (1992) Converting from independent to employee salesforces: The role of perceived switching costs. *J. Marketing Res.* 29(1):101–115.