When to Morph

by

John R. Hauser

Glen L. Urban

and

Guilherme Liberali

August 2010

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

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

Guilherme (Gui) Liberali is a Visiting Scholar at MIT Sloan School of Management, and Assistant Professor of Marketing, Erasmus School of Economics, Erasmus University Rotterdam, 3000 DR Rotterdam, The Netherlands, (1.781) 632-7674, fax (+31)10-408-9169, liberali@mit.edu.

This research was supported by the MIT Sloan School of Management, the Center for Digital Business at MIT (ebusiness.mit.edu), the Suruga Bank, and Applied Marketing Science, Inc. We gratefully acknowledge the contributions of our industrial collaborators, research assistants, and faculty colleagues: Shirley S. Fung, Robert Klein, Clarence Lee, Qiuyuan Jimmy Li, Erin MacDonald, Eleojo Ocholi, Takuya Sugiyama, Yoshio Tokoro, and Kevin Wang.

When to Morph

Abstract

Prior website morphing applications learn between website visitors. They morph early, after a relatively few clicks, so that the final morph dominates consumer response. Morphing early minimizes the costs of switching among morphs. We relax these limitations and address explicitly the tradeoff within a visitor between exploration (waiting to morph in order to get more information about a visitor's cognitive style) and exploitation (morphing early so that the visitor experiences the best morph for as long as possible). We derive a practical dynamic program to optimize when to morph as well as which morph to select. An illustrative example demonstrates the face validity. Reanalysis of a published application suggests improvements in expected sales of at least $5.7, possibly as much as $17 million. Then, in the first field test of a morphing website versus a static website, we test the proposed within-visitor-&-between-visitor algorithm on a Suruga-Bank customer-advocacy website developed for the Japanese card-loan market. The morphing website significantly improves trust, consideration, and purchase probabilities relative to pre-measures. Improvements are larger than obtained by a static website (trust and consideration are significant). The morphing website provides a substantial relative competitive improvement for Suruga Bank in this ¥25 trillion market.

Keywords: *Automated marketing, Bayesian methods, clickstream analysis, cognitive costs, cognitive styles, cultural styles, dynamic programming, field experiments, Internet marketing, optimization, switching costs, trust-based marketing, website design, website morphing.*

1. Exploration versus Exploitation to Determine When to Morph a Website

Website and advertising morphing are diffusing rapidly. Recent applications include:

- BT Group (broadband sales): matching a website's look and feel to visitors' cognitive styles increased projected sales by approximately 20% (Hauser, et al. 2009).
- US Automaker: matching banner advertisements to cognitive styles and the stage of the consumer's buying process increased click-through by over 35%, brand consideration by 24%, brand opinion by 19%, and purchase probability by 7% (personal communication).
- FT Orange: experimenting with matching push advertising to location and purchase interests of mobile device users (Urban, et al. 2009).
- Suruga Bank (card loans): matching a website's look and feel to visitors' cognitive and cultural styles (included in this paper).

In morphing applications an automated system observes the clickstream (or mobile usage) and uses Bayesian methods to update the probabilities that a website visitor or mobile user (hereafter "visitor") is in a particular segment (cognitive style, cultural style, stage in buying process). A Gittins'-index dynamic program then decides which morph to give a visitor by making a (near) optimal tradeoff between learning about the purchase probabilities associated with that morph (exploration) and maximizing sales to the current visitor (exploitation). (If the visitor's cognitive style were known with certainty, a Gittins' solution would be optimal. When cognitive states are partially observable, an expected Gittins' index solution is nearly optimal.) Morphing is computationally feasible because each morph has an index that fully encapsulates the exploration vs. exploitation tradeoff. The firm earns (near) optimal profits by giving the visitor the morph with the largest expected Gittins' index.

All prior applications focused on the decision between visitors. Specifically, websites, banner advertisements, or push advertisements (hereafter, "websites") use Bayesian methods to update beliefs and morph at most once per visitor. For example, the BT Group website morphed after the 10th click. The 10th click was chosen to justify required assumptions that the website morphs sufficiently rapidly so that the final morph dominates the purchase probability and so that the impact of switching costs is negligible.

When the final morph does not dominate, we increase expected profits by solving the early-vs.-later tradeoff within a visitor. If we morph early in a visit, all else equal, the visitor experiences the better morph for a longer period of time. If we morph later in a visit we observe more clicks and can use those clicks to identify better the visitor's segment and target the morph more precisely. To explore this tradeoff we relax simplifying assumptions in Hauser, et al. Specifically, we develop reasonable models of (1) the cognitive costs to changing website characteristics (otherwise we would morph after every click) and (2) how multiple morphs affect visitors' purchase probabilities (not just the final morph).

Using models for switching costs and multiple-morph impacts we develop a dynamic program which identifies the best time to morph (and how often to morph) within a visitor. The dynamic program is optimal within a visitor and optimal asymptotically, but for a finite number of visitors we must address the link to the between-visitor expected Gittins' index solution. We propose and test a linked algorithm. The tests include an illustrative example, a re-analysis of the BT Group website-morphing data, and a field experiment with Suruga Bank. To the best of our knowledge the Suruga Bank experiment is the first test-control field test of website morphing and the first application to try cultural-styles as well as cognitive styles. To set the stage we briefly review prior research.

2. Brief Review of Between-Visitor Website Morphing

2.1 Gittins' Solution Between Respondents

In Hauser et al. website visitors vary by pre-defined segments, r, where each segment represents a particular combination of cognitive styles. A "priming study" develops priors about the probability, p_{rm0}, that a visitor in segment r will make a purchase if he or she is experiences morph m. The Gittins' solution expresses those priors as beta distributions with parameters α_{rm0} and β_{rm0}. If the first visitor from segment r experiences morph m and we observe his or her purchase (or lack thereof), then it is easy to show that the posterior distribution for p_{rm1} is also a beta distribution but with parameters $\alpha_{rm1} = \alpha_{rm0} + \delta_{m1}$ and $\beta_{rm1} = \beta_{rm0} + (1 - \delta_{m1})$ where $\delta_{m1} = 1$ if the first respondent makes a purchase and $\delta_{m1} = 0$ if he or she does not. For the n^{th} visitor we use similar relationships replacing 0 with $n - 1$ and replacing 1 with n.

If we cared only about exploitation we would assign to the n^{th} visitor the morph with the largest expected p_{rmn}. However, suppose another morph m' has a lower $p_{rm'n}$ but a larger variance. It might be optimal to try morph m' for visitor n to learn more about the posterior distribution of $p_{rm'n}$. We can use the knowledge so gained to make better decisions for subsequent visitors. In general, exploration vs. exploitation problems are difficult, but the website morphing problem has a simple but elegant solution. Gittins (1979) proved that the optimal strategy was to derive an index for each potential option (morph) and to choose the morph with the largest "Gittins' index." Gittins' solution is easy to implement because the index, G_{rmn}, for morph m and cognitive-style segment, r, depends only the discount rate, a, and the two parameters of the posterior beta distribution for p_{rmn}. We restate here without derivation the Bellman Equation because we use it later in the paper when linking to within-visitor optimization. This equation is solved iteratively for G_{rmn} and tabled for α and β. (In the Bellman Equation, $R(\alpha_{rmn}, \beta_{rmn}, a)$ is the standard continuation function in dynamic programming.) Discussion, motivation, and derivations are given in Gittins (1979) and, for morphing, in Hauser et al.

(1) $$R(\alpha_{rmn}, \beta_{rmn}, a) = max\left\{\frac{G_{rmn}}{1-a}, \quad \begin{matrix} \frac{\alpha_{rmn}}{\alpha_{rmn}+\beta_{rmn}}[1 + aR(\alpha_{rmn} + 1, \beta_{rmn}, a)] \\ +\frac{\beta_{rmn}}{\alpha_{rmn}+\beta_{rmn}} aR(\alpha_{rmn}, \beta_{rmn} + 1, a) \end{matrix}\right\}$$

As the number of visitors grows ($n \rightarrow \infty$) the variance of the posterior shrinks and the Gittins' index converges to the expected value of the posterior distribution, $G_{rmn} \rightarrow E[p_{rmn}] = \frac{\alpha_{rmn}}{\alpha_{rmn}+\beta_{rmn}}$. We use such asymptotic behavior to simplify the exposition of the within-visitor dynamic program.

2.2 Segment Membership Probabilities

All prior applications (and our application) divide potential visitors into a set of mutually exclusive and collectively exhaustive segments based on cognitive styles (BT Group), cognitive styles and buying-process stage (automotive), or location and purchase interest (FT Orange). Visitors in a each segment differ on their propensities to prefer website characteristics and these differences are reflected in their clickstreams. Let c_{tn} be the t^{th} click observed for the n^{th} visitor. Let $\vec{c}_{tn} = \{c_{1n}, c_{2n}, \dots, c_{tn}\}$ be the vector of clicks up to and including the t^{th} click. Based on a visitor's first t clicks, we derive posterior values for the probabilities, $q_{rn}(\vec{c}_{tn})$, that the n^{th} visitor belongs to the r^{th} segment. Details are in Hauser, et al. For the purposes of this paper we need only that such probabilities can be derived.

2.3 Expected Gittins' Index Solution

While the Gittins' solution is optimal for a known cognitive-style segment, it is not optimal when a visitor's segment is known probabilistically. For partially observable processes the solution is more complicated. Fortunately, Krishnamurthy and Michova (1999) propose and test a simple heuristic that achieves close to 99% of optimality. Their solution is to choose the morph with the highest expected Gittins' index, EG_{mn}, where the expectation is over $q_{rn}(\vec{c}_{tn})$. (The posterior distribution is updated using the $q_{rn}(\vec{c}_{tn})$'s and the δ_{mn}'s.) Hauser, et al. demonstrate that this heuristic performs quite well.

3. A Model of the Visitor's Cognitive Response

To relax the assumption that only the final morph matters, we model switching costs and visitor reaction to multiple morphs.

3.1 Exposure to Multiple Morphs vs. Assuming the Last Morph Dominates

The expected Gittins' index system learns by linking the observed purchase, or lack thereof, to a morph. When visitors see multiple morphs we assign weights, w_t, to clicks to account for the differential impact of early vs. late morphs. 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. Let $\vec{w}$ be the vector of these weights. In equation form, if a visitor sees morph m_t at the t^{th} click, then we modify the expected purchase probability at the end of the visit (prior to switching-cost adjustment) as follows:

$$p_{rn} = \sum_t w_t p_{rm_t} \qquad (2)$$

When convenient we normalize the impact weights so that they sum to 1.0 over clicks.

3.2 Switching Costs

The cost of switching is most salient with an advertising-morphing example. Suppose that after the t^{th} click, posterior probabilities suggest that the n^{th} visitor is in the "committing" stage of his or her search process. The expected Gittins' indices might suggest that the banner advertisement offer a $2,000 rebate on a target automobile. Now suppose that the next few clicks update posteriors such that the visitor is more likely to be in the "comparing" stage. Suppose that the expected Gittins' indices now suggest a competitive test drive. If we morphed to a competitive-test-drive offer from a rebate offer, the visitor would be frustrated when the rebate is withdrawn. A competitive test drive is less effective after the switch from a rebate than it would have been before the visitor was offered a rebate. Switching costs for website morphing are less dramatic, but still real. Even if the visitor prefers an all-text website, he or she may be less enthralled with the look and feel of an all-text website after getting used to a website with many pictures and graphs.

Based on our experience in multiple applications we chose a simple, robust formulation for switching costs. We penalize the expected purchase probabilities by an exogenous application-dependent factor, γ, that is applied each time a switch occurs. We choose a factor rather than a constant increment so that the revised purchase probabilities remain bounded between 0 and 1.

In our application, the impact weights (w_t's) and the switching discount (γ) were set by managerial judgment (Little 1970). The tradeoff between managerial judgment and field measurement is appropriate for the application in this paper. However, with sufficient resources, the impact weights and the switching discount could be set with field experiments. There is nothing conceptually challenging about such experiments; the managerial decision is cost vs. benefit.

3.3 Extensions

Introducing switching costs and impact weights is already a challenging optimization problem and improves upon prior theory and practice, but there are interesting extensions that might be explored in future research. Researchers might make switching costs dependent on the morphs from which and to which the website switches. In addition, switching costs might be modeled as a more complicated function of the number of morphs seen by a visitor. The impact-weight equation might be generalized.

4. Morphing Within a Visit

To simplify exposition we make two simplifications, one generalization, and we temporarily suppress the subscript n that indexes visitors. The first simplification (relaxed later) examines asymptotic behavior—we assume there have been sufficiently many visitors so that we can approximate the Gittins' index with the expected posterior purchase probability ($G_{rm} \rightarrow p_{rm}$). The second simplification assumes that the relative choice of clicks within a morph depends primarily on the characteristics of the clicks. For example, the choice by a visitor between a focused-content <u>link</u> and a general-content <u>link</u> has the same impact on $q_{rn}(\vec{c}_{tn})$ whether the overall morph is graphical or verbal. This assumption makes the decision problem feasible and is reasonable empirically. Finally, for practical purposes we allow more than one click between decisions and thus generalize our notation to allow t to index decision periods rather than clicks.

4.1 When-to-Morph Decision Problem

Figure 1 illustrates the when-to-morph decision problem for the case where the visitor 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 the website displays morph m_t. The respondent makes clicks, c_t, while exploring the website and we update our beliefs about the visitor's cognitive-style segment, $q_r(\vec{c}_t)$. Using the new information, and anticipating more information from subsequent decision periods, we decide which morph, m_{t+1}, to display in the next decision period. To keep track of morph changes, we define $\Delta_{m't}$as an indicator variable such that $\Delta_{m't} = 1$ if we change to morph m' in period t. With this notation, the total number of morph changes is $N_t = N_{t-1} + \sum_{m'} \Delta_{m't}$. To model explicitly the impact of multiple morphs, we represent the purchase decision by δ rather than δ_m as used in prior applications.

Figure 1
When-to-Morph Decision Problem (Within a Visitor, 4-Period Illustration)
(choose morph (m_t) at each decision period to maximize sales, δ)

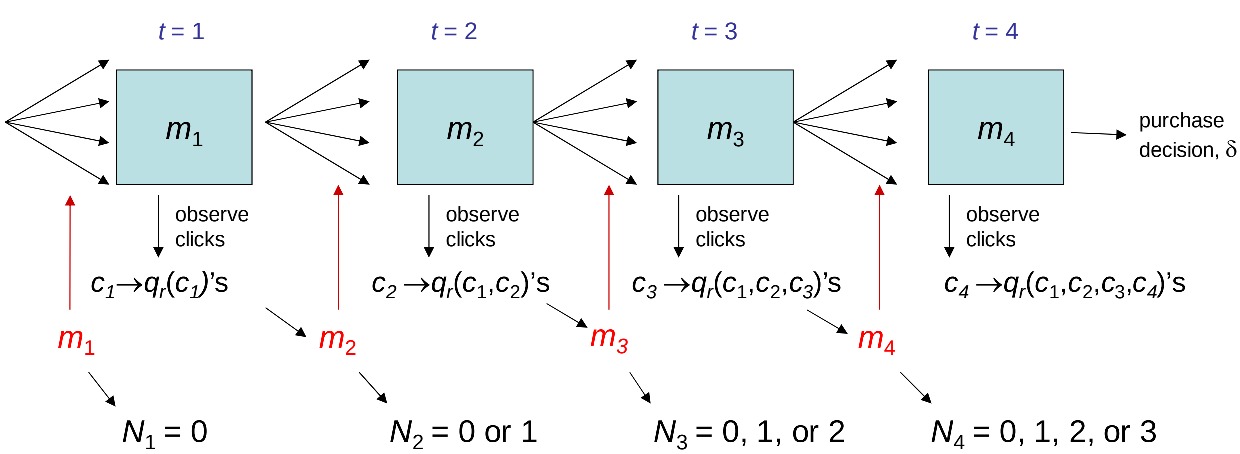

4.2 When-to-Morph Bellman Equation

We solve the when-to-morph dynamic program by deriving the appropriate Bellman equation. To do so we characterize the immediate reward in period t and the value of continuing optimally from period t and beyond.

4.2.1. Immediate reward. Equation 2 implies that rewards are separable over periods thus the immediate reward is based on the impact weight, w_t, the partial effect on the purchase probability of using morph m_t in period t, the number of prior switches, and the terminal probabilities that the visitor belongs to state r, $q_r(\vec{c}_T)$. However, when we make a decision in period t, our best estimate of the terminal probabilities is based on the clicks up to, but not including, the t^{th} decision period. These probabilities $q_r(\vec{c}_{t-1})$ are our expectations over all future clicks. (Implicitly, $c_t, c_{t+1}, \dots, c_T$ are conditioned on $\vec{c}_{t-1}$.) We write the expected immediate reward, $EIR(m_t, \vec{c}_{t-1}, N_{t-1})$, as:

$$\begin{aligned} EIR(m_t, \vec{c}_{t-1}, N_{t-1}) &= \gamma^{N_{t-1}+\Delta_{m_t t}} w_t \sum\nolimits_r E_{c_t, c_{t+1}, \dots, c_T | \vec{c}_{t-1}} q_r(\vec{c}_T) p_{r m_t} \\ &= \gamma^{N_{t-1}+\Delta_{m_t t}} w_t \sum\nolimits_r q_r(\vec{c}_{t-1}) p_{r m_t} \end{aligned} \tag{3}$$

4.2.2. Value of continuing optimally. Although we are uncertain about a visitor's segment, we have expectations about how the future clicks, $c_t, c_{t+1}, c_{t+2}, \dots, c_T$, will evolve based on the visitor's true segment. In turn, our expectations over clicks implies expectations about the evolution of the $q_r(\vec{c}_\tau | \vec{c}_{t-1}, r_{true})$'s for $\tau \geq t$ and for every true cognitive-style segment, r_{true}. These expectations matter for the decision problem. To calculate the expectation of the continuation-value function in the Bellman equation, we track r_{true} in the backwards induction of the dynamic program. When temporarily write the evolution of the clicks as summarized by the evolution of q_r's as conditioned on r_{true}. In notation:

$$V_{t+1}(m_t, \vec{c}_{t-1}, N_t | r_{true}) \equiv E_{c_t, c_{t+1}, \dots, c_T | \vec{c}_{t-1}} V_{t+1}(m_t, \vec{c}_{t-1}, c_t, c_{t+1}, \dots, c_T, N_t \, | r_{true}).$$

But r_{true} is not known when we are making the decision at t, so we keep track of this dependence for $\tau > t$ and take an expectation over this unknown variable at t. That is, when computing $V_\tau(m_{\tau-1}, \vec{c}_{\tau-2}, N_{\tau-1} | r_{true})$ for $\tau > t$, we assume the $q_r(\vec{c}_\tau)$ will evolve as they would if the cognitive-style segment were r_{true}. We also recognize that at the start of the t^{th} decision period, our best estimate of $q_r(\vec{c}_{t-1} | r_{true})$ is $q_r(\vec{c}_{t-1})$, $N_t = N_{t-1} + \Delta_{m_t t}$, and the decision is made based on current expectations of r_{true}. With these substitutions, we have:

$$V_t(m_{t-1}, \vec{c}_{t-1}, N_{t-1}) = \max_{m_t} \left\{ \begin{array}{c} \gamma^{N_{t-1}+\Delta_{m_t t}} w_t \sum\nolimits_r q_r(\vec{c}_{t-1}) p_{r m_t} + \\ \sum\nolimits_s \left[q_s(\vec{c}_{t-1}) \gamma^{\Delta_{m_t t}} V_{t+1}(m_t, \vec{c}_{t-1}, N_{t-1} | r_{true} = s) \right] \end{array} \right\}$$

Finally, we recognize that the multiplicative nature of switching costs enables us to factor out the effect of prior switches, N_{t-1}, when choosing m_t because all future p_{rm_τ}'s are reduced by the same factor, $\gamma^{N_{t-1}}$. However, we keep track of the $\Delta_{m_\tau \tau}$'s for $\tau > t$ when we compute the conditional values, $V_\tau(m_{\tau-1}, \vec{c}_{\tau-2} | r_{true} = s)$. (When computing this conditional continuation value, we use our expectations about the evolution of $q_r(\tau | \vec{c}_{\tau-1}, r_{true})$ for $\tau \geq t$.) We now have the Bellman equation that we use in backward induction.

$$(4) \qquad V_t(m_{t-1}, \vec{c}_{t-1}) = \max_{m_t} \left\{ \gamma^{\Delta_{m_t t}} \left(w_t \sum_r q_r(\vec{c}_{t-1}) p_{rm_t} + \sum_s [q_s(\vec{c}_{t-1}) V_{t+1}(m_t, \vec{c}_{t-1} | s)] \right) \right\}$$

4.2.3. An illustration of the when-to-morph dynamic program. We consider an abstracted problem where the number of decision periods is fixed and the segment probabilities evolve in a known manner dependent only on segment membership. The illustrative problem in Table 1 has four potential morphs and four visitor segments. The first panel gives the morph-dependent purchase probabilities at $t = T$; the second panel gives the expected evolution in the segment probabilities, $q_r(\vec{c}_t | r_{true})$, for $r_{true} = 1$. In this illustration the segment probability evolutions for other true segments are symmetric (and not shown in Table 1); the priors remain the same. In this abstract problem, it is relatively easy to represent Equation 4 in a spreadsheet (available from the authors). Improvements relative to a fixed-time-to-morph website vary from 0% to 50%, depending upon the chosen switching costs, impact weights and other assumptions.

Prior to any observations, the priors and purchase probabilities slightly favor segment $r = 2$ and its corresponding best morph $m = 2$. Not surprisingly, before observing any clicks, the dynamic program solution begins with $m_1 = 2$. For illustration, suppose that the visitor's true segment is $r_{true} = 1$, and suppose that all periods have equal impact on the purchase probabilities. We illustrate the effect of varying switching costs. For this case and with perfect knowledge (and no switching costs), the best morph is $m_t = 1$.

Table 1
Illustrative Morph Assignment Problem for the Within-Visitor Dynamic Program

Table for $p_{rm} = G_{rm}(n \to \infty)$	*r* = 1	*r* = 2	*r* = 3	*r* = 4
$m_T = 1$	0.39	0.24	0.18	0.16
$m_T = 2$	0.17	0.45	0.20	0.15
$m_T = 3$	0.19	0.17	0.34	0.17
$m_T = 4$	0.20	0.15	0.10	0.36

Table for $q_r(\vec{c}_t \mid r_{true} = 1)$		*r* = 1	*r* = 2	*r* = 3	*r* = 4
t = 0	priors	0.24	0.28	0.24	0.24
t = 1	$c_1 \to q_r(\vec{c}_1 \vert r_{true})$	0.40	0.20	0.20	0.20
t = 2	$c_1, c_2 \to q_r(\vec{c}_2 \vert r_{true})$	0.64	0.12	0.12	0.12
t = 3	$c_1, c_2, c_3 \to q_r(\vec{c}_3 \vert r_{true})$	0.91	0.03	0.03	0.03

Without perfect knowledge of r_{true}, the observations, c_1, update the visitor's segment probabilities, $q_r(\vec{c}_1)$. In our illustration the posterior probabilities begin to imply that it is more likely that the visitor's true segment is $r = 1$. When switching costs are low it is optimal to set $m_2 = 1$. On the other hand, when switching costs are more substantial, it is optimal to wait to learn more about the visitor's segment before switching. With greater switching costs the dynamic program continues with $m_2 = 2$. We next observe c_2. The new observations update our believes about segment membership making $r = 1$ more probable. Morph 1 becomes more attractive for this visitor so the dynamic program switches to $m_3 = 1$ in the third period. However, when switching costs are substantially larger (γ smaller) it becomes optimal to wait longer to change morphs. In fact, if switching costs are extremely large, it is never optimal to change morphs. For example, the optimal solution of the illustrative example implies the following optimal paths for the indicated switching costs:

- $\gamma = 0.95 \Rightarrow m_1 = 2, m_2 = 1, m_3 = 1, and\ m_4 = 1$
- $\gamma = 0.80 \Rightarrow m_1 = 2, m_2 = 2, m_3 = 1, and\ m_4 = 1$

- $\gamma = 0.60 \Rightarrow m_1 = 2, m_2 = 2, m_3 = 2, and\ m_4 = 1$
- $\gamma = 0.40 \Rightarrow m_1 = 2, m_2 = 2, m_3 = 2, and\ m_4 = 2$

We can also examine sensitivity to the impact weights. For example, if later periods have larger impact weights it is optimal to wait longer to change morphs. Changing $\vec{w}$ to {0.00, 0.25, 0.25, 0.50} from {0.25, 0.25, 0.25, 0.25} makes it optimal to continue in morph 2 one period longer. If the last period dominates ($\vec{w} =$ {0.00, 0.00, 0.05, 0.95} we approach the assumptions of Hauser, et al. and it is optimal to wait until the dominant last period. Other sensitivity analyses suggest that the optimal solution behaves as we might expect intuitively: with more-rapid learning the dynamic program changes morphs earlier and with increased rewards for continuation beyond period 4 the dynamic program changes morphs later. In summary, the optimal solutions quantify the when-to-morph decision and do so with face validity. Furthermore, the when-to-morph solution reduces to the Hauser-et-al. solution under the assumptions in Hauser, et al.

4.3 When Visitors Leave the Website Earlier or Later

In applications, not all visitors stay for the same number of observation periods. We cannot predict with certainty when they will exit. To model this phenomenon we adjust the Bellman equation to allow the visitor to leave in any period with probability, ψ. For the immediate reward we recalculate effective impact weights w_t' to account recursively for normalizations for the number of periods before exit. For example, before the first period we expect the visitor to continue for one period with probability ψ, for two periods with probability $(1-\psi)\psi$, for three periods with probability $(1-\psi)^2\psi$, and so on.

The continuation reward is ψ times the conditional continuation reward calculated for a visitor exit and $(1-\psi)$ times the conditional continuation reward calculated for a visitor continuing to click. Empirically, the random-exit Bellman equation behaves similarly to the fixed-exit Bellman equation as long as ψ is moderately small. As ψ increases it becomes optimal to morph earlier. (An illustrative spreadsheet is available from the authors for random-exit optimization.)

5. Linking When-to-Morph Optimization with Between-Visitor Gittins' Indices

In the previous section we temporarily assumed asymptotic conditions ($n \to \infty$) such that $G_{rm} \to p_{rm}$. In this case, the between-visitor and the within-visitor dynamic programs are decoupled and Equation 4 provides the optimal solution. For moderate n the two dynamic programs are not independent because each time a visitor experiences a morph we learn a small amount about the posterior distributions of p_{rm}'s. It is a small amount because the purchase observation is distributed over all morphs seen by the visitor and, as in Hauser, et al., distributed over the partially observable cognitive-style segments.

As in Hauser, et al. we seek an heuristic solution that improves upon current practice. However, the solution must run sufficiently fast so that we can morph between clicks during a website visit and so that we maintain the speed of the between-visitor expected-Gittins'-index heuristic. We propose an heuristic which is asymptotically optimal as $n \rightarrow \infty$, improves upon current practice, and runs sufficiently fast. Prior experience in related problems suggests that this heuristic is near optimal for finite n.

5.1 Gittins'-Motivated Solution

We link the within-visitor and between-visitor dynamic programs by using the Gittins' index, G_{rm_t}, rather than the purchase probability p_{rm_t} in the within-visitor dynamic program. This heuristic is optimal asymptotically and equivalent to the Hauser-et-al. algorithm when the final morph dominates. This heuristic is intuitive because Gittins' solutions are often robust and separable. For example, in the "branching bandits" literature researchers have demonstrated that it is optimal to replace the uncertain outcome of an indexable decision process with its Gittins' index (Bertsimas and Niño-Mora 1996; Tsitsiklis 1994; Weber 1992). Our heuristic is not a branching-bandit solution, but shares the property that we are replacing an uncertain decision process with its Gittins' index so that we can solve the within-visitor dynamic program.

Because the initial morph affects outcomes as per Equation 2 we must also select the initial morph. Fortunately, selecting the initial morph is equivalent to the between-visitor problem. Thus, we choose the initial morph with the highest expected Gittins' index where the expectation is over q_{rn} prior to any observed clicks.

5.2 Modified Bayesian Updating

We modify the Bayesian updating of α_{rmn} and β_{rmn} to reflect the distributed nature of the multiple-morph model of consumer response (Equation 2). If we let ζ_{mn} be the impact weights summed over the decision periods that the n^{th} visitor saw morph m, then the posterior values of the Beta parameters are approximated by modifying Equation 2 in Hauser, et al. We correct for the switching discount because we are seeking the posterior distribution for p_{rm}.

$$\alpha_{rmn} = \alpha_{r,m,n-1} + \gamma^{-N_{nT}} q_r(\vec{c}_T)\zeta_{mn}\delta_n \tag{5}$$
$$\beta_{rmn} = \beta_{r,m,n-1} + \gamma^{-N_{nT}} q_r(\vec{c}_T)\zeta_{mn}(1-\delta_n)$$

Following Hauser, et al., Equation 5 uses pseudo-likelihood methods to approximate the posterior distributions. This is necessary for the algorithm to run sufficiently rapidly.

5.3 Summary of the Linked Dynamic Programs

Website morphing combines Bayesian updating and Gittins' indices to allocate website characteristics (morphs) to visitor segments. However, when switching costs matter and morphs other than the final morph affect purchase probabilities, we must address exploration and exploitation within visitors. In the previous section we derived a dynamic program to make this tradeoff optimally when $n \to \infty$. For finite n the within-visitor dynamic program must be linked to the between-visitor Gittins' dynamic program. We propose a simple heuristic that reduces to established website morphing when the final morph dominates and is asymptotically optimal when $n \to \infty$. At minimum we expect the heuristic to improve performance relative to established algorithms. We now test the proposed algorithm by reanalyzing the BT Group data and then test the algorithm empirically via a test-vs.-control field experiment on a Suruga Bank website.

6. Reanalysis of the BT Group Website-Morphing Application

Hauser, et al. applied their between-visitor website-morphing algorithm to a BT Group website that sold broadband service in Great Britain. Their tests were based on purchase probabilities obtained from a priming study of 835 respondents. Eight website morphs varied on three sets of characteristics; visitor segments varied on four sets of ipsative cognitive styles. Hauser, et al. report that website morphing improved sales (profits) which, if implemented system-wide, would represent almost $80 million. The BT website morphed sufficiently early that the purchase probabilities were (mostly) based on the final morph.

To examine the proposed within-visitor-&-between-visitor heuristic we reanalyze their data assuming a switching cost of $\gamma = 0.99$, equal impact weights, and four morphing opportunities. For an apples-to-apples comparison with our empirical test, we select the same number cognitive segments as in the empirical test and simulate 10,000 visitors from each cognitive-style segment (a total of 40,000 synthetic visitors). Visual inspection of G_{rmn} plots indicates that 10,000 synthetic visitors per cognitive-style segment is appropriate to observe performance for finite n and for G_{rmn} convergence. (Relative interpretations are unlikely to change when sixteen segments are simulated. All plots available from the authors.) So that the results are not sensitive to random error, we simulate each optimization method five times and take the average. (The standard deviations are small.)

In Table 2 we compare the performance of the Hauser-et-al. algorithm to the proposed algorithm. We reran the Hauser-et-al. algorithm to account for switching costs which reduce purchase probabilities. In addition, purchase information is now spread among all morphs within a visitor, hence optimization is more difficult for all algorithms—the G_{rm} converge more slowly. As a result, it is natural that the expected rewards and improvements for the Hauser-et-al. algorithm are smaller in Table 2 than in Hauser, et al. (2009, p. 212).

Table 2
Reanalysis of BT Group Website Morphing (with Switching Costs)

	Expected Reward	Improvement [d]
No Gittins' loop nor knowledge of cognitive styles.	0.3165	0.0%
Website morphing (with switching costs)		
Between-visitor morphing only. Hauser-et-al. algorithm. [a]	0.3462	9.4%
Proposed algorithm linking expected Gittins' and within-visitor dynamic programs	0.3497	10.5%
Proposed algorithm focusing on "when to morph" component [b]	0.3586	13.3%
Perfect information on cognitive-styles and on purchase probabilities [c]	0.3828	20.9%

[a] Expected Gittins' algorithm from Hauser, et al. (2009) modified to account for visitor response to multiple morphs and switching costs and applied to the four cognitive-style segments.

[b] Starting morph is chosen as if $G_{rm} \rightarrow p_{rm}$; the within-visitor dynamic program chooses when to morph.

[c] Upper bounds. Applications do not have perfect information on either cognitive styles or purchase probabilities

[d] Improvement relative to that obtainable without any morphing.

We focus first on the decision of when to morph by allowing the algorithm to learn fully the best initial morph. This algorithm improves expected sales by 13.3% compared to the 9.4% improvement obtained by the Hauser-et-al. algorithm. This improvement represents an additional \$17.0 million in expected sales based on the data in Hauser, et al. Next we compare the fully-linked algorithm using the expected Gittins' algorithm to learn the initial morph. The linked algorithm takes longer to learn the initial morph, but still improves expected sales by 10.5%, slightly more than the Hauser-et-al. algorithm. While this difference may seem like a small improvement, it represents $5.7 million in expected sales. (An additional interpretation of Table 2 suggests that the Hauser-et-al. algorithm, which assumes no switching costs and assumes that the final morph dominates, performs quite well even when those assumptions are violated.)

7. Cognitive- and Cultural-Style Morphing at Suruga Bank

In Japan consumers 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 Yen and pays interest when the funds are withdrawn. The terms of card loans vary among banks and can often be confusing. Some banks offer low interest and high limits, but a more-difficult screening process while other banks offer higher interest and lower limits, but an easier screening process. In 2006-2007 Orix spent ¥13.6 billion mostly on banner advertising and Acom spent ¥10.9 billion mostly on television advertising ($\$1 \approx ¥95$).

Suruga Bank is a Japanese commercial bank in the greater Tokyo area. Unlike most commercial banks, it has focused on retail banking for more than twenty 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 ten virtual branches and eight virtual alliances (Tokoro 2008, p. 7). Suruga is less well-known than other Japanese banks, spending approximately $1/10^{th}$ that of Acom and Orix on advertising (¥1.4 million , Tokoro 2008, p. 17). To reach more consumers Suruga developed a consumer advocacy website on which it presented the best products from all competitors without distinction. By using a strategy of openness and honesty Suruga sought to demonstrate that its products (low interest rates, high limits, but careful screening process) would meet the needs of many consumers. As part of Suruga's consumer focus, Suruga's managers experimented with website morphing to match the look and feel of their advocacy site to the varying cognitive and cultural styles of Japanese consumers.

7.1 Cognitive and Cultural Styles, An Evolving Design

Following Hauser, et al., website designers began with a March 2008 priming study of Japanese consumers. E-mail invitations were sent to consumers drawn randomly from a panel of 600,000 consumers maintained by Interface Asia; 5,454 responded. Respondents were offered ¥200 to complete the survey. After screening on age and interest in card loans, 2,114 respondents were invited to visit an experimental website and complete a survey. Of these, 502 respondents

(23.7%) completed the survey, including a requirement to browse the website for at least 2½ minutes and for at least 10 clicks.

There were 16 versions of the experimental website that varied on characteristics such as 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 vs. informal Japanese language, and hierarchical vs. egalitarian images. The website designers sought characteristics that would match the cognitive and cultural styles of Japanese consumers (Allison and Hayes 1996; Hofstede 1983; Riding and Rayner 1998). The website designers sought to implement characteristics that would distinguish the look and feel of the 16 versions along the dimensions of analytic vs. holistic, deliberative vs. impulsive, collectivistic vs. individualistic, and hierarchical vs. egalitarian. The first two ipsative dimensions are targeted to cognitive styles; the second to cultural styles.

Because this was the first Japanese application and the first cultural-style application, Suruga Bank realized that the website designers might not succeed in developing morphs that varied on all four ipsative dimensions. (The website designers were native Japanese.) Thus, in the priming study, we measured visitors' cognitive and cultural styles with 17 scales that have proven to distinguish cognitive and cultural styles in other contexts (see appendix). With these scales we sought to determine empirically how potential customers of Suruga Bank evaluated the website characteristics.

Empirically, the 17 scales suggested that when evaluating the 16 morphs card-loan consumers were best distinguished by two cognitive dimensions rather than combinations of the cognitive and cultural dimensions. However, the morphs that led to the largest stated purchase intentions for each cognitive-style segment varied on cultural styles (as targeted by the website designers). For example, as summarized in Table 3, the best morph for holistic-impulsive consumers was the morph that the web designers believed was holistic, individualistic, hierarchical, and deliberative. This surprise deserves a digression.

Impulsive consumers prefer a website that designers thought would favor deliberative. There are two phenomena at work here. First, website designers are not the same as consumers; website morphing favors the voice of the consumer over the voice of the designer. For example, while website designers thought they were manipulating holistic-vs.-analytic characteristics (columns in Table 3), consumer data suggests that the characteristics had differential influence on impulsive vs. deliberative consumers (rows in Table 3). Second, the characteristics designed for the deliberative morphs may be inherently better—deliberative characteristics are more effective for 3 of the 4 cognitive-style segments. Table 3 illustrates the basic robustness of the website-morphing methodology. By relying on a priming study we identify effective morphs even when the empirical data override designers' beliefs.

While cultural characteristics do not define segments, websites that vary on cultural characteristics influence cognitive-style segments differently. For example, websites designed to be hierarchical tend to be more effective for holistic Japanese consumers and websites designed to be egalitarian tend to be more effective for analytic Japanese consumers. This potential interaction between cultural and cognitive styles is interesting and worth exploration in future research.

Table 3 is sufficient for Suruga Bank and for the purposes of this paper because discrete morphs were identified that affect cognitive-style segments differentially. Given distinct morphs, the "Gittins' engine" identifies the best morph for each segment in the field experiment. Suruga Bank developed a morphing website based on the three morphs (for four segments) in Table 3. In addition, the priming study provided insight that enabled Suruga Bank to identify a fourth morph with high potential. The four morphs were used in the field experiment.

Table 3
Best Morph for Each Cognitive-Style Segment
(Based on the results of a priming study; descriptors are from web-designers' beliefs. *Italics* indicates where observed consumer response overrules designers' beliefs.)

	Holistic Cognitive Style	Analytic Cognitive Style
Impulsive Cognitive Style	hierarchical, individualistic, *deliberative*, holistic	hierarchical, individualistic, *deliberative*, *holistic*
Deliberative Cognitive Style	egalitarian, individualistic, deliberative, *analytic*	egalitarian, collectivistic, *impulsive*, analytic

7.2 Inferring Cognitive Styles from the Clickstream

Website morphing requires that we link the characteristics of visitors' clicks to their cognitive styles. (Following Hauser, et al., we use Bayesian updating based on a matrix, Ω, that summaries the preferences of consumers in each cognitive-style for website characteristics. The Suruga matrix was estimated from the priming-study data.) For Suruga, we identified seven characteristics of clicks, three functional characteristics, six website areas, and a choice of pictures on the home page. Table 4 summarizes the click- and functional-characteristics; Figure 2a illustrates website areas; and Figure 2b illustrates the home page which gave visitors a choice to "click one of the following two pictures to enter the Card Loan Guide." The picture on the left is thought to appeal to deliberative visitors; the picture on the right to impulsive visitors.

Six independent judges, blind to the hypotheses of the research, rated the seven characteristics (reliability = 0.84). The functional characteristics and website areas were binary variables. Because Hauser, et al. (footnote 17) found no statistical difference between Monte Carlo Markov Chain estimation of the posterior means of Ω and maximum likelihood estimation of Ω, we used maximum likelihood estimation for the Suruga application. The model was strongly significant ($p < 0.001$) with a U^2 of 33.9%. There were many significant and intuitive values for the coefficients. For example, impulsive visitors prefer links to fast solutions, advisors, and forums, but not analytic tools or "content directly addressed to you." This intuitive Ω suggests that the Bayesian engine should be able to identify well the cognitive-style segments in the Suruga application. Details of the Ω estimation, scale development, and reliability analyses are available from the authors.

Table 4
Click Characteristics, Functional Characteristics, and Website Areas

Click characteristics	Expect the click to point to a webpage with graphs. Expect the link to point to a webpage with detailed, technical content. Expect the link to point to a webpage with textual content. Expect the link to point to a webpage with options and alternatives. Expect the link to point to a webpage with content regarding popular trends. Expect the link to point to a webpage with content directly addressed to you. Expect the link to point to a webpage which uses formal language.
Functional Characteristics	Provides information Analytical tool Graphical elements

Figure 2
Opening Webpage for Suruga Bank Card Loan Website

8. Implementation of Website Morphing at Suruga Bank

Suruga's experimental website matched the four morphs to visitors using the within-visitor-&-between-visitor "when-to-morph" algorithm. Suruga sought to test the website with 1,000+ visitors—the first test-vs.-control experiment comparing a morphing website to a static website.

With more resources, we would like to evaluate both the concept of website morphing and the incremental improvement due to the when-to-morph algorithm. However, 1,000+ visitors is not many. Website morphing is designed for high-traffic websites. For example, the simulations in §6 were based on 10,000 visitors per cognitive-style segment. The BT simulations suggest that the Gittins' indices settle down between the 250th and the 350th visitor, but continue substantial exploration through the 1,000th visitor (Hauser, et al., Figure 3). Based on our simulations, we estimated that the Suruga website would begin to settle down around the 1,000th visitor. Because of this resource constraint we were able to test only one algorithm. We chose the best algorithm, the when-to-morph algorithm. Suruga allocated additional visitors to a static website, approximately one-third of the visitors that they allocated to the test website.

The Suruga application is a proof of concept for the when-to-morph algorithm and a test-vs.-control experiment comparing that algorithm to a static website. Future field experiments can parse the incremental value of the within-visitor portion of the when-to-morph algorithm.

In November-December 2009 Suruga recruited consumers from the Interface Asia panel. E-mail invitations were sent to 62,000 potential respondents; 13,696 responded (22.1%). Screening and incentives were similar to those in the priming study. On net, 3,514 consumers were directed to the card-loan site, of which 1,997 explored the website for at least 2½ minutes and 10 clicks (56.9%). Of these consumers, 1,400 completed pre- and post-visit questionnaires providing data with which to evaluate the websites (70.1%). Of these, 1,066 experienced the morphing website and 334 experienced the static website. To avoid an obvious demand artifact, the website was not identified as a Suruga Bank website—recall that the website provides competitive information on Acom, Orix, and Suruga.

8.1 Results of the Field Experiment

We report results for three dependent measures: trust, consideration, and purchase probability. (Due to human-subjects considerations, we were not allowed to actually sell card loans on the websites.) Purchase probability is self-explanatory and is the long-term goal of Suruga Bank. However, because Suruga Bank is new to card loans it is important to Suruga Bank that the website improve consideration. Trust is also central to Suruga Bank's strategic initiatives. Not only does increased trust build a relationship with customers, but prior research suggests that the effect of competitive information on consideration and purchase is mediated through trust (Bart, et al. 2005; Liberali, Urban and Hauser 2010). Methodologically, trust is a more-continuous measure (5-point scale) and likely to be more sensitive to changes than the discrete measure of consideration (0 vs. 1). (Purchase probabilities are conditional measures among those who consider Suruga Bank.) Greater sensitivity is important given the resource constraint of approximately 1,000+ visitors.

Table 5 summarizes the results. Suruga's website, whether static or morphing, increased consideration and purchase intentions significantly and Suruga's morphing website increased trust significantly ($p < .01$ for all comparisons). The only anomaly is that the static website seems to have decreased trust ($p = .06$). The more interesting comparison is between the test vs. control improvements. The improvement in trust, consideration, and purchase probabilities are larger for the morphing website than for the static website. The difference in trust is significant ($p < .01$), the difference in consideration is marginally significant ($p = .08$), and the difference in purchase intention probabilities is not significant ($p = .16$). Taken together these measures suggest that the morphing website outperformed the static website, at least with respect to trust and consideration and possibility with respect to purchase probabilities. And these results are based on a sample of just 1,000+ visitors when the Gittins' indices are still settling down.[1] We would have likely obtained even better results had we been able to observe enough visitors for the Gittins' indices to converge to their asymptotic values ($G_{rm} \rightarrow p_{rm}$).

Table 5
Results of the Suruga Bank Field Experiment

	Pre-measures		Post-measures		Pre-Post Difference	
	Static Website	Morphing Website	Static Website	Morphing Website	Static Website	Morphing Website
Trust (5-point scale)	2.19	2.36	2.05	2.48 [a]	-0.14	0.12 [b]
Consideration	6.3%	4.7%	26.9% [a]	30.1% [a]	20.7%	25.4% [c]
Purchase	2.4%	1.9%	6.9% [a]	8.6% [a]	4.5%	6.6% [d]
Sample size	334	1,066	334	1,066	334	1,066

[a] Post-measures significantly larger than pre-measures at <0.01 level.
[b] Morphing website significantly larger than static website at < 0.01 level.
[c] Morphing website significantly larger than static website at 0.08 level.
[d] Morphing website larger, but not significantly larger, than static website (p = 0.16).

8.2 Suruga's Competitive Advantage is Enhanced by Website Morphing

Suruga Bank is currently a smaller competitor in card loans. If the a morphing website enhances trust, consideration, and purchase probabilities for Suruga, the profit potential is significant. We re-examine Table 5 to evaluate whether the morphing website provides a competitive advantage for Suruga Bank.

The morphing website leads to substantial improvements in consideration and purchase probabilities (642% and 443%, respectively) . However, these are likely inflated by forced exposure to the website and a general demand artifact that might have increased consideration and purchase probabilities for all banks. (Recall that the website was not identified as a Suruga website.) Acom and Orix also increased by 327% and 203%, respectively. Using the Acom and Orix

[1] Although we have pre- and post-measures on just 1,066 visitors, the Gittins' indices were being updated through the 1,531 visitors who were assigned to the morphing website (3/4[ths] of the 1,997 total visitors).

improvements to normalize the Suruga improvements suggests that the morphing website provides approximately a two-fold relative increase in consideration and purchase probabilities for Suruga compared to Acom and Orix.

However, the static website was also a customer advocacy website that provided competitive comparisons. We observed substantial improvements in consideration and purchase probabilities due to the static website (429% and 285% for Suruga, 403% and 329% for Acom/Orix). Performing the same calculations we estimate that the static website provides approximately the same improvement for Suruga as for Acom and Orix. In other words, the morphing website provides a better competitive advantage to Suruga Bank than does the static website. In this proof-of-concept test visitors could not actually apply for card loans and we do not know if they would have been approved—approval depends upon their credit ratings, data which we did not collect due to privacy concerns. Nonetheless, considering that the 2008 card-loan market in Japan accounted for over ¥25 trillion in total available balances ($265 billion, personal communication from Suruga Bank), this potential two-fold relative competitive advantage for Suruga Bank is clearly worth further exploration.

9. Summary and Future Directions

In this paper we extended the Hauser-et-al. website morphing methodology to model cognitive switching costs and to address the impact of multiple morphs during a website visit. We developed a dynamic program which, in steady state, chooses the optimal time to morph during the website visit. We propose and test an algorithm to link the within-visitor when-to-morph dynamic program to the between-visitor expected-Gittins'-index dynamic program. The linked algorithm automatically identifies which morph to provide a visitor, when to provide that morph, and how many times to morph.

We tested the proposed algorithm in three ways. First, we illustrated the when-to-morph dynamic program with a focused example to show that the optimal solution has face validity. Second, we demonstrated that the algorithm improves performance using the data from a previously published application. This test establishes that the proposed algorithm outperforms existing methods. Third, we compare an experimental morphing website to a static website in a field experiment with Suruga Bank. The morphing website provides a significant increase in trust, a marginally significant increase in consideration, and an increase in purchase probabilities. The competitive implications of these increases for Suruga Bank are substantial.

The Suruga proof-of-concept test is the first field test of website morphing and the first attempt to include cultural as well as cognitive styles as a basis for morphing. We believe that the improvements and field test in this paper demonstrate the growing potential of morphing.

There are still many challenges. (1) The switching-cost and impact parameters were set by managerial judgment in the Suruga application. Future experiments could provide empirical values for these parameters. (2) The sample-size constraints limited the Suruga comparison to a morphing website vs. a static website. Future experiments might parse the incremental value of within-visitor morphing. (3) While the proposed algorithm improves on current practice, its optimality depends upon asymptotic n. We believe it is close to optimal for finite n, but a formal demonstration is beyond the scope of this paper. (4) Cultural styles played a role in the design of the Suruga website, but visitor segments defined by cultural styles were not justified empirically. The literature is rich in the study of cultural styles; we expect future applications to improve upon the Suruga experience. (5) Ongoing applications include morphing banner advertisements, push advertising, and mobile-phone apps. For example, Suruga Bank is in the process of developing and testing a mobile-phone morphing website. Within-visitor morphing should provide improvements to those applications.

References

Allinson, Christopher W. and John Hayes (1996), "The Cognitive Style Index: A Measure of Intuition-Analysis for Organizational Research," *Journal of Management Studies*, 33, 1, (January), 119-135.

Bart, Yakov, Venkatesh Shankar, Fareena Sultan, and Glen L. Urban (2005), "Are the Drivers and Role of Online Trust the Same for All Web Sites and Consumers? A Large-Scale Exploratory Empirical Study," *Journal of Marketing*, 69, (October), 133-152.

Bertsimas, Dimitris and José Niňo-Mora (1996), "Conservation Laws, Extended Polymatroids And Multi-Armed Bandit Problems," *Mathematics of Operations Research*, 21, 2, (May), 257-306.

Frederick, Shane (2005), "Cognitive Reflection and Decision Making." *Journal of Economic Perspectives*, 19 (4). 25-42.

Gittins, John C. (1979), "Bandit Processes and Dynamic Allocation Indices," *Journal of the Royal Statistical Society, Series B (Methodological)*, 41, 2, 148-177, plus commentary.

Hauser, John R., Glen L. Urban, Guilherme Liberali, and Michael Braun (2009), "Website Morphing," *Marketing Science*, 28, 2, (March-April), 202-224.

Hofstede, Geert (1983), "National Cultures in Four Dimensions," *International Studies of Management and Organization*,13, 46-74.

Krishnamurthy, Vikram and Josipa Mickova (1999), "Finite Dimensional Algorithms For The Hidden Markov Model Multi-Armed Bandit Problem," *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 5, 2865-2868.

Liberali, Guilherme, Glen L. Urban, and John R. Hauser (2010), " Does Providing Competitive Information to Your Own Customers Increase Sales?," Working paper, (Cambridge, MA: MIT Sloan School of Management).

Little, John D. C. (1970), "Managers and Models: The Concept of a Decision Calculus," *Management Science*, 16 (8) B466-485.

Riding, Richard J. and Stephen Rayner (1998), *Cognitive Styles and Learning Strategies: Understanding Style Differences in Learning and Behavior*, (London, UK: David Fulton Publishers).

Tokoro, Yoshio (2008), "Improving Customer Experience through Advocacy and Morphing: A Web Application for Suruga Bank," S. M. Thesis, (Cambridge, MA :MIT Sloan School of Management).

Tsitsiklis, John N. (1994), "A Short Proof of the Gittins Index Theorem," *Annals of Applied Probability*, 4, 1, (February), 194-199.

Weber, Richard (1992), "On The Gittins Index For Multi-Armed Bandits," *Annals of Applied Probability*, 2, 4, (November), 1024-1033

Urban, Glen L., John R. Hauser, Guilherme Liberali, Michael Braun, and Fareena Sultan (2009), "Morph the Web to Build Empathy, Trust, and Sales," *Sloan Management Review*, 50, 4, (Summer), 53-61.

Appendix
Cognitive and Cultural Scales (Expected Prior Grouping)

Analytic vs. holistic	I see what I read in mental pictures. I will read an explanation of a graphic/chart before I try to understand the graph/chart on my own. I enjoy deciphering graphs, charts, and diagrams about products and services. I like detailed explanations. I'm usually more interested in parts and details than in the whole. I am detail-oriented and start with the details in order to build a complete picture.
Deliberative vs. impulsive	I prefer planning before acting. I like to make purchases without thinking too much about the consequences. A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost? [impulsive = 10 cents] [a] If it takes 5 machines to produce 5 widgets. How long does it take 100 machines to produce 100 widgets? [impulsive = 100 minutes] [a] In a lake, there is a patch of lily pads. Every day the patch doubles in size. If it takes 48 days for the patch to cover the lake, how long would it take for the patch to cover half of the lake? [impulsive = 24 days] [a]
Collectivistic vs. individualistic	In choosing my ideal job it would be important to have sufficient time for my personal life. I buy products in order to differentiate myself from other people. Buying products for my family and friends is more important to me than buying things for myself.
Hierarchical vs. egalitarian	I think authority and leadership are very important in life. I'm usually afraid to express disagreement with my superiors or with important persons. I value mostly experts' opinions when I buy a product.

[a] Modified from Frederick's (2005) Cognitive Reflection Index