2007

Simple Models of Discrete Choice and Their Performance in Bandit Experiments

Noah Gans
University of Pennsylvania

George Knox

Rachel Croson
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/oid_papers

Part of the [Business Administration, Management, and Operations Commons](#), and the [Organization Development Commons](#)

Recommended Citation

Gans, N., Knox, G., & Croson, R. (2007). Simple Models of Discrete Choice and Their Performance in Bandit Experiments. *Manufacturing & Service Operations Management*, 9 (4), 383-408. <http://dx.doi.org/10.1287/msom.1060.0130>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/135
For more information, please contact repository@pobox.upenn.edu.

Simple Models of Discrete Choice and Their Performance in Bandit Experiments

Abstract

Recent operations management papers model customers as solving multiarmed bandit problems, positing that consumers use a particular heuristic when choosing among suppliers. These papers then analyze the resulting competition among suppliers and mathematically characterize the equilibrium actions. There remains a question, however, as to whether the original customer models on which the analyses are built are reasonable representations of actual consumer choice. In this paper, we empirically investigate how well these choice rules match actual performance as people solve two-armed Bernoulli bandit problems. We find that some of the most analytically tractable models perform best in tests of model fit. We also find that the expected number of consecutive trials of a given supplier is increasing in its expected quality level, with increasing differences, a result consistent with the models' predictions as well as with loyalty effects described in the popular management literature.

Keywords

service quality, multiarmed bandit, decision-making under uncertainty, discrete-choice models, experimental tests

Disciplines

Business Administration, Management, and Operations | Organization Development

Simple Models of Discrete Choice and Their Performance in Bandit Experiments *

Noah Gans George Knox Rachel Croson
OPIM Department Marketing Department OPIM Department

The Wharton School, University of Pennsylvania
Philadelphia, PA 19104

`{gans,gknox,crosonr}@wharton.upenn.edu`

December 16, 2004

Abstract

Recent operations management papers model customers as solving multi-armed bandit problems, positing that consumers use a particular heuristic when choosing among suppliers. These papers then analyze the resulting competition among suppliers and mathematically characterize the equilibrium actions. There remains a question, however, as to whether the original customer models upon which the analyses are built are reasonable representations of actual consumer choice.

In this paper, we empirically investigate how well these choice rules match actual performance as people solve two-armed Bernoulli bandit problems. We find that some of the most analytically tractable models perform best in tests of model fit. We also find that the expected number of consecutive trials of a given supplier is increasing and convex its expected quality level, a result that is consistent with the models' predictions, as well as with loyalty effects described in the popular management literature.

*Thanks to Gabriel Silvasi for developing the software for the experiments, as well as for providing computing support for the analysis of the data. Thanks also to Eliza Ponsz, who facilitated the testing of the software and recruiting of subjects. We are grateful to Lawrence Brown, Dean Foster, Wes Hutchinson, and Bob Meyer for suggestions concerning the analysis. We also thank seminar participants at the University of Minnesota, the Yale School of Management, and the University of Pennsylvania's Decision Processes Brown Bag Series for their helpful comments. The research of Rachel Croson was supported by NSF Grant SBR-9876079, which provided funds to pay the subjects who participated in the experiments. The research of Noah Gans was supported by NSF Grant SBR-9733739, The Wharton Financial Institutions Center, and The Fishman-Davidson Center for Service and Operations Management.

1 Introduction

In many business contexts, customers switch among suppliers. This switching is often observed in service industries, where it is called customer “defection” or “churn.” It is also a common phenomenon in the context of consumer products, where brand switching is a widely studied phenomenon. Among the factors responsible for this switching is random variation in the utility or value that a customer obtains each time he or she patronizes a supplier.

More specifically, in many settings, the quality of service offered to customers has an inherently random component. Competing suppliers’ quality distributions, in turn, jointly determine customer switching behavior and market shares. By installing extra capacity or instituting additional quality control measures, a service provider can improve the distribution of its service quality, at a cost.

The mechanism by which customer switching occurs is naturally modelled as the solution to a multi-armed bandit problem. While the bandit problem is a straightforward means of representing a customer’s repeated choices, it is difficult to work with analytically, and this difficulty raises two sets of questions. Customers faced with bandit problems need practical characterizations of optimal (or near-optimal) behavior so that they make better decisions. Companies, which are the “arms” in consumers’ bandit problems, need to better understand of how people actually make choices in a bandit setting so they can effectively respond to their customers’ needs.

A number of recent operations management papers have sought to recommend actions to companies that are arms in bandit or bandit-like problems (e.g., Gans (2002a), Gaur and Park (2003), and Hall and Porteus (2000)). These papers begin by positing that customers use a particular heuristic for making repeated choices under uncertainty. They then derive aggregate statistics regarding customers’ choices, such as each supplier’s “customer share,” as a function of a given set of supplier quality distributions.¹ Finally, the papers model suppliers as competitively choosing quality distributions, and they characterize suppliers’ resulting equilibrium choices.

A common feature of these papers is the analytical tractability of their choice models. That is, the rules customers are assumed follow are not complex, and this allows the resulting expressions for aggregate choice statistics to be simple functions of firms’ quality distributions.² While this simplicity facilitates the papers’ competitive analyses, there remains a question as to whether or not the underlying customer models are reasonable representations of actual consumer choice.

In this paper, we report the results of one effort at validating these models of consumer choice.

¹Customer share refers to the fraction of all purchases that a customer makes at a given supplier.

²It is worth noting that a model’s analytical tractability can be viewed in at least two ways. On the one hand, one may view a model which requires less computation or memory on the part of subjects as being more tractable. On the other, one may view as more tractable a model for which aggregate measures of a subject’s performance are simple functions of system parameters, preferably with simple, closed-form mathematical expressions. We are primarily concerned with the latter, and when we write “tractable” this is what we mean. Nevertheless it is typically the case that the two meanings go hand in hand.

We view the models we test as representing a variety of tradeoffs between analytical tractability and richness of the representation of customers’ learning processes. To the extent that more stylized, tractable models adequately capture customer choice behavior, they may be valuably used in the type of competitive analysis described above. Alternatively, if more complex models, whose statistics are more difficult to derive, are significantly better predictors customer behavior, then the use of highly stylized models needs to be rethought.

We consider two sets of models. One set includes three models that represent successive approximations of a normative analysis of the problem. The first of these is a Gittins-index model, which posits that subjects make choices that maximize the expected discounted value of their outcomes and use Bayes’s rule to incorporate learning from past experience. The next is a Myopic analogue to the Gittins-index rule, in which subjects are Bayesian but choose to maximize the expected value of only their immediate choice. The last model, which we call Simple, further reduces the model of myopic behavior by assuming that subjects categorize arms as being either “good” or “bad.”

The second set includes three related models of choice that are outgrowths of the literature on statistical learning and decision-making under uncertainty. The first of these is a version of the myopic rule, described above, in which subjects remember the outcome of only the last n trials. A further simplification is a Hot Hand rule, which represents subjects as staying with the current arm until it loses in n consecutive trials. The last is an Exponential Smoothing model in which subjects update their beliefs about the quality of an arm by taking a weighted average of their past beliefs with the current outcome. In Section 2 we review the literature related to the proposed models and their tests, and in Section 3 we formally define the two sets of models to be analyzed.

We use experiments to test the performance of the models, and we analyze the results in two ways. First, we consider a prediction that is consistent across the representations: the expected switching time (the expected number of trials that a subject consecutively samples from a given arm before switching to an alternative) should be increasing and convex in the expected payout of the arm. The increasing-convex property is also consistent with the claim, made in the popular management literature, that marginal increases in service quality can have increasingly dramatic benefits in customer loyalty and lifetime value (Jones and Sasser 1995).

Second, we distinguish among the various models’ ability to predict subjects’ behavior, choice by choice. The use of experiments to distinguish among models in a horse race like this is common (Roth 1988). The experimental setting must be carefully constructed so that all models make clear predictions and there exists a dimension along which they can be evaluated. Our experiments, described in Section 4, are designed to do just this.

Our findings, presented in Sections 5 and 6, are as follows. First, measures of the expected switching time are increasing and convex in the expected payout of an arm. Second, more analytically tractable models provide the best fit to subjects’ observed choices. In contrast, the more complex Gittins-index model, which maximizes expected discounted rewards, fares most poorly in tests of model fit.

Our experimental results also suggest that there is no single model that dominates the others. That is, there exist significant subgroups of subjects for whom different choice models perform best. This finding suggests that suppliers' choices of service-quality distribution(s) should explicitly account for various market segments, each segment characterized by a different model of choice. These conclusions and other extensions are discussed in Section 7.

2 Literature Review

Bandit-like models have been used in a wide variety of contexts. Classic examples include clinical drug trials and petroleum exploration (e.g., Gittins (1989), Banks et al. (1997), Anderson (2001)), and in consumer product or service settings (e.g., Meyer and Shi (1995) and Erdem and Keane (1996)). For a more complete description of the multi-armed bandit problem, please see Appendix A.

Because of its importance in many different fields, normative aspects of the bandit problem have received considerable attention over the years, and we briefly describe previous work devoted to settings like ours, an infinite-horizon, discounted version of the problem. Perhaps most widely known is the work of Gittins and Jones (1974) who proved the following: from each arm's Markov chain and current state an index may be calculated, and at any stage it is optimal to choose the arm with the highest index. Gittins (1979) went on to develop explicit expressions for this so-called Gittins index, though in the context of Bayesian bandits (with arms that are members of exponential families of distributions) the index has proven to be tremendously difficult to calculate. In turn, Chang and Lai (1987) developed approximations to the Gittins index for Bayesian bandits, which they proved to be asymptotically optimal as the discount rate approaches one, so that the decision-maker samples infinitely before switching away from an arm. (See also Brezzi and Lai (2002).)

The difficulty in calculating optimal choices has also led to suggestions for simpler, heuristic solutions to the bandit problem. An early paper by Schmalensee (1975) analyzes the use of the Bush and Mosteller (1955) linear learning model for solving the problem. In the computer science literature, various other forms of reinforcement learning have been applied to the problem as well (Sutton and Barto 1998).

This difficulty naturally leads one to wonder how well people actually solve bandit problems. Surprisingly, however, there is relatively little positive work that investigates how people make choices in bandit settings. In particular, we are aware of only four such papers. Horowitz (1973), Meyer and Shi (1995) and Banks et al. (1997) use Bernoulli bandits, whose arms have simple win-lose payout distributions. Anderson (2001) uses arms with more complex payout distributions, such as simulated dice rolls and normally-distributed rewards.

Results from these studies suggest that people's behavior is roughly consistent with predictions from the Bayesian model, but with important and systematic deviations. Horowitz (1973) and Meyer and Shi (1995) offer exploratory analyses, using experimental data to generate hypotheses concerning

the heuristics used by subjects. Both papers find that subjects’ choices reflect a Bayesian updating of priors, but also a bias toward choices that are more myopic than are optimal. Horowitz (1973) also finds that subjects over-sample from inferior arms (over-experiment), a bias which is distinct from that of myopic behavior. Banks et al. (1997) and Anderson (2001) devise experiments that provide sharp tests of specific hypotheses concerning the presence of certain heuristics or biases in choice behavior. Anderson (2001) finds that subjects experiment less than optimally and provides evidence that risk aversion associated with more diffuse priors is the likely cause of the effect. Banks et al. (1997) manipulate the arms’ prior distributions so that, in some cases, optimal choice behavior is myopic, and in others it is not, and they find that subjects’ behavior is consistent with these normative predictions.

The intent of our analysis differs from the work described above. Rather than using the data either in an inductive fashion, to propose new models of choice, or to construct sharp tests of hypotheses concerning specific biases away from expected utility-maximizing behavior, we take an “engineer’s” view of the problem. Our aim is to identify models of choice behavior that can easily render aggregate statistics that are useful in the context of competitive analysis (such as expected switching time or fraction of times chosen). Thus, we judge a model’s value along two dimensions: its analytical tractability (for operations management professionals), as well as its ability to represent the choices that a wide variety of people make.³

Our approach is similar in spirit to recent work in the economics literature by Harless and Camerer (1994) and Camerer and Ho (1999). In the former, the authors consider generalizations of expected utility theory along two dimensions: “fit” with the observed data, as well as “parsimony” (parallel to our notion of tractability). In the latter, the authors apply tools from marketing research to judge the fit of models to subjects’ choices in experiments, much as one would fit brand choice models to panel data. We will pursue a similar strategy here.

The models we test come from a variety of sources. Two (Myopic and Simple) are simplifications of the original Gittins-index model. They are motivated by two sets of findings from empirical work on perception and decision-making. The first observation is that people tend to choose more myopically than is optimal (e.g., Horowitz (1973) and Meyer and Shi (1995)). The Myopic model explicitly incorporates this bias. The second is that people appear to systematically categorize as they make sense of their experiences. That is, they maintain mental examples of how entities in the world behave, and they interpret an experience with an entity by comparing their perceptions with the typical or exemplary characteristics of their mental picture of how that category behaves. This structure shows up in Kahneman and Tversky’s well-known “representativeness heuristic” (Kahneman and Tversky 1973, Tversky and Kahneman 1974), and it forms the basis of category and exemplar theory in

³Hutchinson and Meyer (1994) suggest that a formal positive theory of sequential choice is likely include combinations of simple rules or strategies, each of which optimizes a limited or restricted version of the task at hand. In this context, one may view our analysis as explicitly considering the efficacy of various simple policies, while postponing the larger issue of how people use the rules in concert.

cognitive psychology (e.g., Henderson and Peterson (1992)). The Simple model assumes that subjects categorize arms as “good” or “bad” and that subjects exhibit the same myopia that is present in the Myopic model.

Models in the second set that we test are variants of those with origins in behavioral decision theory. Our “Last- n ” model is based on the finding that bounds on memory limit subjects’ ability to recall long histories of previous outcomes (Miller 1956). The “Hot Hand” model, described in Gilovich et al. (1985), posits that subjects erroneously ascribe positive serial correlation to the outcomes of repeated trials. (The mirror of this is the negative correlation implied in “gambler’s fallacy” rules (e.g., Burns and Corpus (2004)). The classic, exponential smoothing model, which was first developed in the early 1960’s (Brown and Meyer 1961), was introduced into the marketing literature in Guadagni and Little (1983), and more recently proposed as a choice model in March (1996).

3 The Bandit Problem and Choice Models

In this section we formally define the multi-armed bandit problem addressed in this paper, as well as the models of choice that we will test. For more detail on the bandit problem, please see Appendix A.

3.1 The Multi-Armed Bandit

The multi-armed bandit problem used in our study is defined as follows. There exist m arms, indexed $i = 1, \dots, m$. When sampled, arm i yields a randomly distributed reward (or value or utility) with fixed distribution, U^i . Formally, we describe U^i as uniquely defined by a parameter $\theta \in \Theta$ and as having cumulative distribution function $F(u|\theta)$, so that $U^i \sim F(u|\theta^i)$. Let $\mu_\theta = E[U|\theta] = \int u dF(u|\theta)$.

A subject must repeatedly choose among the m arms, and we $t = 1, 2, \dots$ be the time index of her choices. While she knows that each arm has a fixed distribution $\theta^i \in \Theta$, she does not know what the various θ^i ’s are. Her task is to repeatedly choose among the arms in a way that maximizes the aggregate value of her choices.

Let a *policy* $\pi = \{\pi(1), \pi(2) \dots\}$ be a sequence of choices of arms, and let Π be the class of policies that is nonanticipating with respect to future rewards. Then we formally define the subject’s problem as that of finding (and executing) a policy, $\pi \in \Pi$, that will maximize the expected discounted value of the future stream of rewards, $\sup_{\pi \in \Pi} E_\pi [\sum_{t=1}^{\infty} \alpha^t U_t]$, where $\alpha \in (0, 1)$ is the one-period discount rate.

3.2 The Gittins Index for the Bayesian Bandit and Related Models

A Bayesian subject may view each arm’s θ^i as a random variable with support on Θ . For arm i she maintains a cumulative distribution function $P_t^i(\theta)$ that represents her understanding at time t of the

quality distribution, $\theta^i \in \Theta$, under which she believes the arm to be operating. Let $\{P_0^1, \dots, P_0^m\}$ be the “prior” information the subject has before she begins her sequence of choices.

After each choice, the Bayesian subject uses the new sample, U_t , and Bayes’s rule to update her beliefs. Specifically, if she uses arm i at time t and receives reward u , then her new (posterior) belief distribution will be

$$dP_t^i(\theta|u) = \frac{dP_{t-1}^i(\theta)dF(u|\theta)}{\int_{\Theta} dP_{t-1}^i(\theta)dF(u|\theta)} \quad \forall \theta \in \Theta. \quad (1)$$

If she does not sample from i , then $dP_t^i(\theta) = dP_{t-1}^i(\theta) \quad \forall \theta \in \Theta$.

Gittins Index Model

The so-called Gittins index of arm j describes the expected discounted reward per unit of expected discounted time,

$$G(P_t^j) \triangleq \sup_{\mathcal{T} > t} \left\{ \frac{\mathbb{E} \left[\mathbb{E} \left[\sum_{s=t+1}^{\mathcal{T}-1} \alpha^{s-t} U(P_{s-1}^j) | P_t^j \right] \right]}{\mathbb{E} \left[\mathbb{E} \left[\sum_{s=t+1}^{\mathcal{T}-1} \alpha^{s-t} | P_t^j \right] \right]} \right\}, \quad (2)$$

where $\mathcal{T}$ is a stopping time with respect to the history of the process through time $t-1$. The notation, $U(P_{s-1}^j)$, emphasizes the fact that the marginal (subjective) distribution of value at $(s-1)$ is a function of the distribution of the subject’s belief at the time. To maximize her expected discounted reward, a subject should choose the arm with the largest Gittins index, $i = \arg \max_j \{G(P_{t-1}^j)\}$ (Gittins and Jones 1974, Gittins 1979).

Myopic Model

In practice, $\mathcal{T}$ and the Gittins index are extremely difficult to compute, and experimental evidence suggests that people behave more myopically than is optimal (Horowitz 1973, Meyer and Shi 1995). In contrast, a *myopic* customer significantly simplifies the determination of her preferred arm by ignoring the option of future switching. For each arm j she more simply calculates the expected discounted reward, given she remains on arm j for all time,

$$M(P_t^j) \triangleq \frac{\sum_{s=t+1}^{\infty} \alpha^{s-t} \mathbb{E}[U(P_{s-1}^j)]}{\sum_{s=t+1}^{\infty} \alpha^{s-t}} = \mathbb{E}[U(P_t^j)] = \int \mu_{\theta} dP_t^j(\theta), \quad (3)$$

She then chooses the arm i that maximizes this long-run expected reward: $i = \arg \max_j \{M(P_t^j)\}$. From the right hand side of (3), we see that this is equivalent to choosing the arm that myopically maximizes her expected reward at t .

Like her “rational” counterpart, the myopic subject uses the reward realized at i and Bayes’s rule (1) to calculate her posterior beliefs, P_{t+1}^i , concerning i ’s quality distribution. Again, for all $j \neq i$, $P_{t+1}^j = P_t^j$.

Simple Model

While the Myopic subject’s task is significantly simpler than that of her Gittins-index counterpart, it

still may be quite complex. She must maintain a prior distribution over the set, Θ . She must also perform potentially difficult integrations to update the prior distributions (1) and to calculate the indices (3).

Indeed, a common finding in behavioral research is that people may substitute simpler heuristics for these complex, integrative tasks. Perhaps the best-known example of this type of simplifications is Kahneman and Tversky's "representative heuristic" (Kahneman and Tversky 1973, Tversky and Kahneman 1974). The *Simple* model uses categorization to further simplify the integrative aspects of the Myopic model.

In addition to using (3) to myopically choose the arm that maximizes immediate expected reward, a Simple subject partitions the arms' possible quality levels into two categories – good and bad – with respective reward distributions $F_G \equiv F(u|\theta^G)$ and $F_B \equiv F(u|\theta^B)$. That is, she further simplifies the choice process of the Myopic subject by reducing the set of distributions that she recognizes to be $\Theta = \{\theta^B, \theta^G\}$. In turn, P_t^i , the prior distribution she maintains for each arm's θ^i , collapses to be p_t^i , the probability that i is good rather than bad. Thus, rather than judging *how* good or bad an arm is, the subject's problem is more simply to decide *whether* an arm is good or bad.

Let $\mu_G \triangleq E[U|\theta^G]$ and $\mu_B \triangleq E[U|\theta^B]$. Using (3) we can define the index used by the Simple subject in terms that are exactly analogous to $M(P_t^i)$:

$$S(p_t^i) = E[U(p_t^i)] = \mu_G p_t^i + \mu_B (1 - p_t^i). \quad (4)$$

Again, $S(p_t^i)$ is the expected reward of sampling from i at t . Then, given a realization, u , the Simple subject's use of Bayes's rule (1) to update the prior probability that i is good reduces to

$$p_t^i = \frac{p_{t-1}^i dF_G(u)}{p_{t-1}^i dF_G(u) + (1 - p_{t-1}^i) dF_B(u)} = \left(1 + \frac{1 - p_{t-1}^i}{p_{t-1}^i} \cdot \frac{dF_B(u)}{dF_G(u)} \right)^{-1}. \quad (5)$$

Algebraic manipulation shows that the index defined by (4)–(5) is equivalent to

$$\tilde{S}_{t-1}^i = X_0 + \sum_{s=1}^{t-1} 1\{\pi(s) = i\} \cdot X_s, \quad (6)$$

where $X_0 = \ln\left(\frac{p_0^i}{1-p_0^i}\right)$ is a log-likelihood that reflects the subject's initial belief concerning the probability that arm i is good (p_0^i), $X_s = \ln\left(\frac{dF_G(U_s)}{dF_B(U_s)}\right)$ is the log-likelihood that the s th trial comes from a good arm, and $1\{\cdot\}$ is the indicator function. That is, $S(p_{t-1}^i) > S(p_{t-1}^j)$ if and only if $\tilde{S}_{t-1}^i > \tilde{S}_{t-1}^j$.

Note that (6) is a random walk which has an intuitive interpretation. In it, X_0 is an initial level of satisfaction that a subject has for arm i . Each time a subject samples from i , her experience leads an immediate response to the quality of the interaction, X_s . In turn, this immediate response is integrated into the subject's overall satisfaction with i in a straightforward, additive fashion: better outcomes increase and worse outcomes decrease overall satisfaction.

Conversely, one can view the Simple model as being defined as (6), an additive model of “satisfaction” in which the arm with the highest cumulative satisfaction is chosen in each trial. This is precisely the model of “cumulative discrete choice” proposed in Gilboa and Pazgal (2001). Given this second definition, we then note that the Simple model is also consistent with the behavior of a myopic Bayesian who categorizes arms as good or bad.

The Simple model is also consistent with the behavior of a so-called “satisficing” subject. That is, rather than seeking the arm that maximizes expected rewards, the Simple subject will be satisfied with any arm that meets some target level of average reward per trial, μ^* . If $\mathbb{E}[U^i] < \mu^*$, then arm i does not meet the subject’s required satisfaction level, and the expected number of times she will sample from i before switching is finite. If, however, $\mathbb{E}[U^i] \geq \mu^*$, then arm i meets the subject’s so-called “aspiration level,” and she is expected to continue sampling from i indefinitely. The twin notions of satisficing and aspiration levels have a long history, dating back to Simon (1959) and beyond.

The Simple model was used in Gans (2002a) in the context of models of service competition, and the paper provides a general, closed-form, representation of μ^* for exponential families of probability distributions. In Section 6, we also provide an explicit representation μ^* in the context of our experimental setting.

The models described above form a hierarchical family. Most complex is the Gittins-index model, which derives from the rational behavior of Bayesian subjects. Less complex is the Myopic model, and then least complex is the Simple model, which may be viewed as a categorical version of the Myopic model, as well as an additive, “random walk” index of a subject’s satisfaction with a given arm.

3.3 Other Choice Models

There exist many other possible representations of choice under uncertainty that can be tested in our Bandit setting. In this paper, we concentrate on three: a “Last- n ” model, which is analogous to a Myopic model with limited memory of past trials; an “Exponential Smoothing” analogue to the Simple model; and a “Hot Hand” model that reacts to recent wins and losses on the currently sampled arm.

Last- n

It is well-known that individuals have only limited memory. For example, a long stream of research in psychology documents that individuals can remember roughly 7 pieces of information, such as digits of telephone numbers, after which they need record-keeping or other external aids (Miller 1956).

The *Last- n* index explicitly incorporates the effect of limited memory, using only the results of the last n trials on a given arm. Formally, it is calculated in the same fashion as the Myopic index, the difference being that the Last- n index for an arm corresponds to a Myopic index in which only the previous n trials on that arm are remembered.

Although the model represents subjects as having limited memory of past events, in a significant sense its use is more demanding of subjects' memory than the related Myopic rule. More specifically, let $k_i(t) = \sum_{s=1}^t 1\{\pi(s) = i\}$ be the number of times a subject has sampled arm i by time t , and let $s_i(k) = \min\{t \mid k_i(t) = k\}$ be the time of the k^{th} sample from arm i . Then a subject that uses the Last- n model must always recall each of the last n samples on each arm; that is, given she has sampled $j \geq n$ times from arm i , she must recall samples $\{U_{s_i(j-n+1)}, \dots, U_{s_i(j)}\}$. In contrast, to update of the analogous Myopic index a subject need only recall the prior distribution, P_{t-1}^i , along with the outcome at time t , U_t .

Hot Hand

A subject that uses the Hot Hand (HH) model ascribes positive serial correlation to the trials of a repeated random sample. The rule's underlying premise is that an arm that has recently won is more likely (than average) to win again.⁴

We define a "HH- n " rule in our experiments as follows. If the subject experiences n consecutive losses in the last n trials on an arm she should switch to the other arm; otherwise she should continue to sample from the current arm. Thus, another way of stating a HH-1 rule is "stick on a winner and switch on a loser." Similarly, a HH- n rule could be described as switch only on n consecutive losers. Like the Last- n rule, the HH- n rule requires that a subject maintain a detailed record of the outcome of each of the last n trials on a given arm.

The HH- n family of rules differs significantly from the other rules that we test in two fundamental ways. First, as far as we can tell, it is only directly applicable to Bernoulli outcomes. While outcomes with more complex distributions can be reduced to Bernoulli trials by applying a threshold – so that outcomes above the threshold are considered to be wins and those below the threshold, losses – all of the other rules immediately generalize to arbitrarily complex distributions, without having to be transformed in this manner. Second, the HH- n family is not an index rule. All of the other rules that we test calculate an index for each arm and recommend that the subject sample from the arm with the higher index. In contrast, in the HH- n rule the decision to stay on or switch away from the current arm is based only on that arm's performance. Information concerning the past performance of alternative arms is not used at all in the switching decision.

We note that Hall and Porteus (2000) use a variant of the HH-1 rule in their analysis of service competition. In this model, a customer that receives satisfactory service remains with the current supplier, while a customer that receives unsatisfactory service switches to a competitor with a fixed probability p .

⁴The name "hot hand" derives from basketball, in which there is a common belief that players have "hot" hands, so that their success probabilities in making free-throw attempts exhibit positive serial correlation. In the context of problems of repeated choice under uncertainty, we say a belief in the "hot hand" implies that a subject (erroneously) believes that a future payout is (positively) serially correlated with recent performance (Gilovich et al. 1985).

Exponential Smoothing

Exponential smoothing is a weighted analogue of the Simple model’s additive random walk. Formally, we define the Exponential Smoothing (ES) model as follows. At $t = 0$ the subject’s prior estimate of the average reward to be gained by sampling from arm i is $ES_i(0)$. Then at each trial t , at which the subject samples from arm i , we let

$$ES_i(t) = \gamma U_t + (1 - \gamma) ES_i(t - 1) \quad (7)$$

where $0 < \gamma < 1$. Again, for all $j \neq i$, $ES_j(t) = ES_j(t - 1)$.

In contrast to the index for the Simple model, $\tilde{S}_i(t)$, for which each previous trial carries an equal weight, $ES_i(t)$ is more strongly affected by recent trials at i . To see this, again let $s_i(k)$ denote the time index of the k^{th} trial at arm i . Then, given arm i has been pulled j times by time t , we can write $ES_i(t) = \sum_{k=0}^j \gamma (1 - \gamma)^{j-k} U_{s_i(k)}$, where $U_{s_i(0)} \equiv \gamma^{-1} ES_i(0)$. Thus, for a fixed $0 < \gamma < 1$, the weight of the $s_i(k)^{th}$ trial at i declines geometrically (roughly exponentially) quickly with each new sample from i , hence the name exponential smoothing. The higher the weighting factor, γ , the more quickly the weight declines. For $\gamma = 0$, the index remains constant, $ES_i(t) = ES_i(0)$ for all t . Conversely, for $\gamma = 1$, a subject’s index is defined by her most recent trial: $ES_i(t) = \max_{\{s \leq t | \pi(s)=i\}} U_s$.

Exponential smoothing models have a long history of use in a number of fields related to learning. An early reference from the forecasting literature is Brown and Meyer (1961). Guadagni and Little (1983) is a well-known application from the marketing literature, which applies smoothing to a bandit-like brand choice problem. A recent example from the learning literature in psychology is March (1996), which uses simulation to analyze properties of smoothing models.

Of particular interest to us is a generalization of a smoothing model that appears in Gaur and Park (2003), a paper that analyzes service-level competition among inventory systems. Here outcomes are Bernoulli – either a customer order is filled or it is not – and the model is extended so that the smoothing constant, γ , may differ, depending on whether or not an order is filled. (This also echoes asymmetric elements of the stochastic learning models that date back to Bush and Mosteller (1955).)

It is also worth noting that the smoothing model is consistent with the belief that an arm’s reward distribution is Markov modulated, rather than i.i.d. across trials. In this case, it can be shown that smoothing is analogous to the use of a Bayesian procedure in which prior trials’ results are discounted. (For details, see Matsuda and Sekiguchi (1971).) In the context of our Bandit problems, one may interpret a subject’s use of exponential smoothing as “not believing” that arms are i.i.d. and, therefore, discounting earlier sample information, since it is more likely to have been obtained from a reward distribution that differs from the one that is currently being sampled.

Additional Models

There is a potentially vast array of additional models that one might also consider. For example, we have not analyzed models of reinforcement learning that can be found in the computer-science

and economics literatures (e.g., Sutton and Barto (1998) and Erev and Roth (1998)). Nor have we considered the use of various combinations of the current models.

Nevertheless, there does exist one model which we *have* analyzed but do not include in the body of the paper. This is the so-called “probability-matching” model that has a long history within the research literature on human and animal choice (e.g., Robbins and Warner (1973)).

The reason we do not explicitly include the results is two-fold. First, the model is not directly applicable to the bandit setting. Second, when a natural variant of the model (which is applicable to the setting) is fit to subject data, it performs quite poorly. For more on the model and results, see Appendix H.⁵

To better understand how well these models apply to the bandit problem, we conduct two sets of tests. The first seeks to validate or refute a general prediction that is consistent across all models: that the expected number of consecutive trials on an arm is increasing and convex in the average quality (reward) of the arm. The second is a more detailed discrimination among the various models, based on the choices that they recommend for a given observed history. Before we provide the details of the tests, we describe the experimental setup.

4 Experimental Setup and Preliminary Results

To empirically test the models of customer response, we have developed and run experiments in which subjects face the bandit problems described above. The experiments are carefully designed to allow us to perform the more general test for convexity, as well as the more detailed, trial-by-trial analysis of each model’s consistency with subjects’ observed choices.

In this section we describe, in detail, how the experiments are structured. We also provide some descriptive statistics that give a general sense of the subjects’ performance in the experiments.

The Value Distribution

The models we test are, ultimately, meant to represent customer switching due to variation in product or service quality. The hedonic nature of these experiences is difficult to control or monitor, however. For example, the same physical measures of service quality – such as speed, accuracy, or courtesy – may be perceived or valued differently by different people.

Therefore, we operationalize differences in perceived value with money; that is, the dollar reward received when sampling an arm is a proxy for the value received. This payment procedure is standard practice in experimental economics and is intended to induce participants to take their task seriously, since real money is at stake (Friedman and Sunder 1994).

⁵To interpret the appendix’s Figure 20, please first read Section 6.

We define the payout distributions of the arms to be Bernoulli random variables: with probability $P\{i \text{ wins}\}$ arm i pays \$0.10, and with probability $P\{i \text{ loses}\} = 1 - P\{i \text{ wins}\}$ it pays nothing. Note that the shifting or scaling of outcomes $\{\$0.00, \$0.10\}$ does not affect the choices the models recommend.⁶ Furthermore, the use of two-point payout distributions allows us to avoid problems associated with subject utilities that may vary nonlinearly with outcomes, controlling for risk preferences that vary across subjects.

One concern about the magnitude of the payments is the steepness or flatness of the resulting curve of total rewards obtained by subjects (Harrison 1989). In addressing this issue we considered a spectrum of alternatives, from running experiments that were hypothetical (as in previous research), in which the reward function is everywhere zero, to those in which individuals make very few, high-stakes choices. Previous research has demonstrated that the move from hypothetical decisions (with no payment) to real decisions (with small payments) produces a significant change in subject behavior, while the move from small payments to large payments does not significantly affect behavior (Camerer and Hogarth 1999).

Thus, our payment scheme has sought to balance the need for payoffs that are responsive to subject choices with that for collecting enough data to adequately distinguish among models. The \$0.10 per win reflects a small, but significant reward for making careful choices, and it has allowed us to collect ample data to fit our models (347 choices per subject). Furthermore, a number of descriptive statistics (reported below) suggest that most participants understood the experiment and paid attention to the task at hand.

The Number and Nature of Arms

Every subject plays a series of three two-armed Bernoulli bandit problems. We will sometimes refer to each problem as one of three “sessions” in which a subject participates. Participants are informed that, in a given session, each of the two arms has a fixed, but unknown, probability of winning; that is, the probability of success may vary from arm to arm, but will remain constant over time for an individual arm.

Prior Information

Before they begin, participants do not know the two arms’ probabilities of winning. We do, however, show participants some prior information about the arms. Specifically, we report (in writing) that each of the two arms has been sampled from three times and that two of the three trials were successes. Given this information, subjects know that both arms have $P\{\text{win}\} \in (0, 1)$, so they are not degenerate. Subjects also know that the prior performance of the two arms is equivalent.

⁶Theoretically, we can normalize the reward from winning to one and that from losing to zero. The expected reward from choosing an arm is the probability of winning on that arm.

Discounting

Recall that the form of the Gittins-index result (2) for rational subjects depends on an infinite horizon setting with constant discounting. In experiments, a common method operationalizing a constant discount rate is to generate a constant probability $(1 - \alpha)$ that any given round of the experiment will be the last. In our experiment $\alpha = 0.99$, and our instructions communicated to subjects that there would be a 1% chance that any given trial would be the last (equivalently, a 99% chance that it would not be the last).

To determine the number of trials to be run, we then implemented the randomized procedure only once, *before* any of the experiments were run, and we used the same three sample outcomes for all subjects: 95, 117, and 135 rounds. In this fashion, the procedure ensured that results would be directly comparable across participants.

Experimental Treatments

The experimental design involved three treatments, one for each of the three bandit problems each subject encountered. Treatments were run within-subject, counterbalanced. Thus, while each participant saw all three treatments, roughly 1/6 of the participants saw each of the six possible orderings of the treatments. The treatments compared arms with probabilities of winning of 0.15 and 0.40, 0.40 and 0.40, and 0.65 and 0.40. This design allows us to test hypotheses about convexity by comparing the aggregate frequency of choices of an arm, given its win-rate. We can also test predictions at the individual level, since each participant saw all three pairs of arms.

Software-Based Implementation

The experiment was implemented via computer on a web page which could be accessed by a typical browser. Participants were told the web site’s address, logged into the system, and played the game.

The bandit problems were implemented as the repeated choice between two colored decks of cards. In each of the three sessions the success probabilities (0.15/0.40, 0.40/0.40, 0.40/0.65) and colors (red/blue, green/gray; yellow/purple) assigned to the two decks and their location on the screen (left and right) were randomly selected for each participant.

The decks were composed of cards that state “YOU WIN!” or “YOU LOSE!” and their compositions remained constant over time. Each time a subject chose one of the available decks, an animation played that showed the deck being shuffled, one card being chosen at random, and the outcome of the trial, win or lose. The card was then replaced, and the deck reverted to its initial state.

At all times subjects also saw the balance of their winnings as it accumulated at \$0.10 per win. Participants could click a “history” button, which would display summary statistics concerning total wins and losses, as well as the entire history of their choices and the resulting outcomes.

At the end of each of the three bandit problems, each subject answered a short questionnaire.

Demographic and other information was collected once the entire game had ended. Appendix B presents the entire set of instructions and participant views of the experiment.

All participants earned a \$5 participation fee plus their accumulated earnings from the experiment. At the end of the experiment, subjects were informed of their total earnings and asked to print a receipt from their browsers. They brought the receipt to an assistant, who checked it against their ID, obtained their signature and paid them their earnings.

Data Collected

Three-hundred and seventy-three participants logged onto the system. Of these, 227 completed the experiment. Thus we have collected data on 347 (95+117+135) choices times 227 participants = 78,769 choices between two options.

For each of the 227 participants, we have a complete record of each of three treatments of the bandit problem, as well as answers to end-of-treatment and end-of-experiment questions. An example of one session of a subject is shown in Figure 1.

Descriptive Statistics

Of the 227 subjects that completed the experiment, 32 had at least one session in which the last “run” – a set of consecutive trials on a given arm – is the only run on that arm. In these cases, the sole run on that arm is artificially truncated by the end of the experiment, and tests for the convexity of the average run length become difficult to perform. We have therefore excluded these subjects from our tests of convexity and of model fit, and our final data set includes 195 subjects, 347 trials per subject, for a grand total of 67,665 recorded choices.

The 195 subjects took, on average, 12 minutes and 53 seconds to complete all three sessions, or about 2.2 seconds per decision. As they moved from the first to the last session, subjects also took less time to make each decision. This may be due either to learning of the experimental setup or to boredom or fatigue. Subjects’ payoffs from the sessions ranged between \$17.50 and \$23.50, with an average of \$20.50.

A number of statistics suggest that the variation in total rewards was sufficient to motivate participants to understand the experiment and pay attention to the task at hand. Only 20 of 195 subjects reported being confused by some aspect of the experiment. During the experiment, 56% of the subjects viewed the history screen (which summarized their previous choices and the outcomes) at some time, and those that viewed the history visited this screen an average of 9.4 times during the exercise. After each session, we also asked subjects which arm has the higher probability of reward: left, right, or neither. Overall, subjects answered correctly more than 69% of the time; and when the arms had different $P\{\text{win}\}$ s, more than 80% answered correctly.⁷

⁷When $P\{\text{left wins}\} = P\{\text{right wins}\} = 0.40$, the fraction of correct answers – that the two arms had the same probability of reward – dropped to 37%. Of course, to say that the two arms had the same $P\{\text{win}\}$ is a point prediction

	Left	Right		Left	Right		Left	Right
1	LOSE		46	WIN		91	LOSE	
2		WIN	47	LOSE		92		LOSE
3		WIN	48		WIN	93	WIN	
4		WIN	49		WIN	94		WIN
5	WIN		50		WIN	95		WIN
6		WIN	51		WIN	96		WIN
7	WIN		52		LOSE	97		WIN
8	LOSE		53	WIN		98		WIN
9		WIN	54		WIN	99		WIN
10		WIN	55		LOSE	100		LOSE
11		WIN	56	LOSE		101	WIN	
12		LOSE	57		LOSE	102		LOSE
13	LOSE		58		WIN	103	LOSE	
14		WIN	59		LOSE	104		WIN
15		LOSE	60	LOSE		105		WIN
16	WIN		61		LOSE	106		WIN
17		LOSE	62		WIN	107		WIN
18		WIN	63		WIN	108		WIN
19		LOSE	64		LOSE	109		WIN
20	LOSE		65		WIN	110		WIN
21		LOSE	66		WIN	111		LOSE
22		WIN	67		WIN	112		LOSE
23		WIN	68		WIN	113	LOSE	
24		WIN	69		WIN	114		WIN
25		WIN	70		WIN	115		WIN
26		WIN	71		LOSE	116		LOSE
27		WIN	72		LOSE	117		WIN
28		WIN	73	WIN		118	WIN	
29		LOSE	74		LOSE	119		WIN
30		LOSE	75	WIN		120		LOSE
31	WIN		76	LOSE		121	WIN	
32	LOSE		77		WIN	122		WIN
33		WIN	78		LOSE	123		WIN
34		WIN	79	WIN		124		LOSE
35		WIN	80		WIN	125		WIN
36		WIN	81	LOSE		126		WIN
37		WIN	82		LOSE	127		WIN
38		WIN	83		WIN	128		LOSE
39		WIN	84		WIN	129	WIN	
40		WIN	85		LOSE	130		LOSE
41		LOSE	86	LOSE		131	WIN	
42	LOSE		87		WIN	132		WIN
43		WIN	88		WIN	133	LOSE	
44		LOSE	89		LOSE	134		WIN
45		LOSE	90		LOSE	135		WIN

Figure 1: Results for Subject 32, Session 2. $P\{\text{left wins}\} = 0.40$ and $P\{\text{right wins}\} = 0.65$.

We also considered how summary statistics varied across experimental conditions. For each subject, we calculated the average run length in a given session, the total number of trials (95, 117, or 135) divided by the total number of runs in that session, and compared average run lengths across the that is a less likely outcome than “left greater than right” or “right greater than left.”

different arm conditions. When both arm probabilities were equal (0.40/0.40) the average run length was the smallest, at 7.6. Average run length was 8.6 when the choice was 0.15/0.40 and 13.3 when the choice was 0.40/0.65. On average, subjects also took slightly more time – about 0.1 seconds more per round – to complete the sessions when the two arms’ probabilities were the same (0.40/0.40), rather than different (0.15/0.40 or 0.40/0.65). Thus, when the arms’ winning probabilities were the same, so there was no clear best arm, there was more switching and subjects spent more time per trial.

Longer run lengths and times per arm were also positively correlated with total earnings. Ordinary least squares (OLS) linear regression of subjects’ earnings on average run length (defined as $\bar{\tau}$ in §5) and on average time per trial showed that a one-unit increase in average run length was associated with a \$0.07 average increase in total winnings, and a one-minute increase in total completion time (equivalently, a 0.173 second increase in time per round) was associated with a \$0.03 average increase in total winnings.⁸ Of course, the fact that larger run lengths were associated with higher earnings may reflect the fact that subjects tended to stay on “better” arms for longer runs.⁹

5 The Expected Switching Time

Suppose that at some arbitrary period, t , a subject has last sampled from arm i , and let τ be the number of additional periods that she will continue to sample from i , before switching to a competitor. Then asymptotic results suggest that the expected switching time, $E[\tau]$ is increasing and convex in the average quality (reward) provided by i . (See Appendix C.)

A test of this property is interesting to us for two reasons. First, it is a basic check of whether or not subjects’ actual behavior corresponds to a prediction that is common to many models. Second, if true the increasing-convex property is also consistent with a broader claim which is made in the popular management literature: that marginal increases in service quality can have increasingly dramatic benefits in customer loyalty and lifetime value (e.g., Jones and Sasser (1995)). Therefore, our first set of tests focuses on $E[\tau]$.

5.1 Estimation

In constructing tests for monotonicity and convexity, we control the quality of the arm and measure the resulting switching times. In the context of Bernoulli arms, $P\{\text{win}\}$ is the measure of quality, and in the experiments, every subject plays three pairs of arms with the same probabilities of winning:

⁸ $R^2 = 0.15$ in the regression, and both coefficients were significant at the 0.05 level.

⁹We have also calculated descriptive statistics for the subjects who were excluded from our main analyses. The 32 subjects who completed the experiment (but had a single, truncated run on a given arm) spent about 4 minutes less total time (about 0.7 seconds less per decision) on average and won approximately the same amount as other subjects. Subjects who did not finish the experiment spent about 0.2 seconds more time per decision and won about \$0.0015 less per decision than the 195 subjects included in our analyses.

0.15 versus 0.4; 0.4 versus 0.4; and 0.65 versus 0.4. The quality points we have chosen are evenly spaced, with $0.65 - 0.40 = 0.40 - 0.15 = 0.25$, to facilitate testing for convexity. Furthermore, in each treatment at least one of the arms has $P\{\text{win}\} = 0.40$, so that the quality of the arm that is *not* used in the convexity test remains constant across treatments. In the 0.40 versus 0.40 treatment, either arm may be used in the convexity tests, and for this treatment we report the results of both arms.

The estimate of $E[\tau]$ is more difficult to calculate. One easily calculated measure is the average number of consecutive trials – or average “run length” – on each arm, calculated as

$$\bar{\tau}_i = \frac{\text{total number of trials on an arm}}{\text{total number of runs on an arm}}, \quad (8)$$

where $i = l$ for the left arm and $i = r$ for the right.

We test that $\bar{\tau}$ is increasing and convex in $P\{\text{win}\}$ as follows. Let $\bar{\tau}_{p,j}$ be the average run length for an arm with $P\{\text{win}\} = p$ and subject j . For each subject we first calculate the first difference between adjacent quality pairs: $\Delta_{1,j}(\bar{\tau}) = \bar{\tau}_{0.40,j} - \bar{\tau}_{0.15,j}$ and $\Delta_{2,j}(\bar{\tau}) = \bar{\tau}_{0.65,j} - \bar{\tau}_{0.40,j}$. Given the experimental setup, we know that these sample differences are independent across subjects, and we use the Wilcoxon Signed Rank test to check that the median of the differences is greater than zero (Lehmann 1975). Our test for convexity runs along the same lines. Here we define each subject’s second difference as $\Delta_j^2(\bar{\tau}) = \Delta_{2,j}(\bar{\tau}) - \Delta_{1,j}(\bar{\tau})$, and we check that it is positive.

Note that there exist two potential problems associated with using $\bar{\tau}_i$ as an estimate of $E[\tau]$. First, subjects’ last runs are censored. For example, from Figure 1 we see that Subject 32’s last run consists of trials 134 and 135 on the right arm. At this point, the experimental treatment was ended. Had the subject been allowed to continue sampling, the run length might have been longer than two, however. A second potential problem concerns the possibility that the sequence of run lengths is not stationary. For example, the sequence of run lengths may be (stochastically) increasing or decreasing, rather than stationary. Indeed, our results suggest that over the course of a session there is a mild increase in run lengths.¹⁰

The overall impact of these effects is not immediately clear to us. On the one hand, the length of the censored run is longer than what was recorded. In this sense, censoring biases the average downward. On the other hand, there is an inspection bias present: runs that are censored are likely to be longer than average. Even if run lengths are not stationary, to the extent that τ is *stochastically* increasing the quality of the arm, then these sample averages should be increasing in quality as well. Nevertheless, there are modifications that can correct for the problem.

One simple alternative, which we denote $\hat{\tau}$, eliminates the last, truncated run from the calculation of the sample average run length. As with the $\bar{\tau}_{p,j}$ ’s, we can use subjects’ $\hat{\tau}_{p,j}$ ’s to calculate $\Delta_{1,j}(\hat{\tau})$, $\Delta_{2,j}(\hat{\tau})$, and $\Delta_j^2(\hat{\tau})$ and use the Wilcoxon Signed Rank test to check for the increasing and convex properties.

¹⁰On average, the number of switches decreased from 12.65 in the first half of each session to 9.97 in the second half, a reduction of about 21%.

A second alternative tests for convexity of a *single* run, starting in a given trial of the session, and we measure the length of the run that begins at trial one, which we call τ_1 . By choosing the first round, we ensure that runs are not truncated. (The only subjects whose first runs are truncated are those that never change arms.) The choice of the first round also ensures that all subjects have exactly the same information about the arms as the run begins, so that differences among run lengths do not result from experience or informational differences.

In many respects, the second alternative is preferable. It has the drawback, however, of not allowing for the within-subject comparisons across the three arms' qualities. This is because, given our randomization scheme, the success probability of the arm first chosen by a subject is not controllable. When faced with 0.15/0.4 or 0.4/0.65 treatments, many subjects (unknowingly) first chose the arm that had a probability of winning of 0.40.

5.2 Results

While we have calculated relevant statistics for all three measures of run length, Figure 2 displays confidence intervals only for $\bar{\tau}$ and τ_1 . Because the results for $\hat{\tau}$ are similar to those for $\bar{\tau}$ we omit their graphical display.

Figure 2: Convexity Results for the Average Run Length

The figure's left panel shows confidence intervals for $\bar{\tau}_p = \frac{1}{n} \sum_{j=1}^n \bar{\tau}_{p,j}$, $p \in \{0.15, 0.40, 0.65\}$. Each of the intervals shown in the panel is calculated as a sample average ± 2 times the standard error of the estimate of the mean. In each of the three intervals, the reported arm competes against an arm with $P\{\text{win}\} = 0.40$, and in the case of 0.40 versus 0.40 treatment, we report the results of both arms.

Here, the sample averages are clearly increasing and convex, and the confidence intervals do not overlap, a further indication that the increasing property holds. Wilcoxon Signed Rank Tests for first and second differences are all vanishingly small (reported as 0 in S-Plus).

For $\hat{\tau}$, the confidence intervals are qualitatively the same as for $\bar{\tau}$. (See Appendix D.) Similarly

the p -values of the Wilcoxon Signed Rank tests were also negligible. Thus, the results support the hypothesis that $\bar{\tau}$ and $\hat{\tau}$ are both increasing and convex in $P\{\text{win}\}$.

The right panel of Figure 2 shows the results for the single run that starts at the first trial. Here, we do not have paired data across the three treatment conditions. Furthermore, the numbers of subjects for which we have results varies across the conditions: for $P\{\text{win}\} = 0.15$ the number of subjects is $n = 106$; for 0.40, $n = 379$; and for 0.65, $n = 100$. In this case, the point estimate of the average length appears to be convex, though *not* increasing from a $P\{\text{win}\}$ of 0.15 to 0.40. Furthermore, the relatively smaller sample sizes for $P\{\text{win}\}$ of 0.15 and 0.65 result in confidence intervals that are much wider than that for $P\{\text{win}\}$ of 0.40.

Although the lack of pairing of the data makes it more difficult to test for convexity, we can use a Mann-Whitney test (also known as the Wilcoxon Rank Sum test; see Lehmann (1975)) to check whether or not the medians are increasing from one treatment condition to the next. While τ_1 is not significantly increasing from 0.15 to 0.40, it is increasing from 0.40 to 0.65, as well as over the whole range, 0.15 to 0.65.¹¹

In summary, the results of this section are generally consistent with the hypothesis that $E[\tau]$ is increasing and convex in the quality of an arm. While each estimate has particular weaknesses, taken together they confirm that this fundamental prediction is consistent with subjects' actual behavior.

6 Estimating the Models' Fit to the Data

In this section, we consider the more difficult question of distinguishing among the various models' fits to the data. We begin the section by describing how the models of Section 3 are applied to the experimental data and how we calculate the indices associated with each of the models. We then describe our estimation procedures in fitting subjects' choices. Finally, we compare how well the models fit subjects' actual choices.

Our results show that the Exponential Smoothing and Hot Hand models provide the best in-sample fits, both on an aggregate and subject-by-subject basis. We also show, however, that the Hot Hand model's fit is largely an artifact of subjects' long run lengths. Among the hierarchy of Gittins-index derived models, the Simple model performed best. Thus there exist models that provide *both* analytical tractability and a reasonable fit to subjects' observed choices.

¹¹The p -values were as follows: 0.34 for τ_1 increasing as $P\{\text{win}\}$ increases from 0.15 to 0.40; and 0.02 for τ_1 increasing as $P\{\text{win}\}$ increases from 0.40 to 0.65 and from 0.15 to 0.65.

6.1 Calculating the Models’ Indices

We first describe how the models’ indices are calculated in the context of the experiment’s Bernoulli outcomes.

Gittins-Index Model

For the Gittins-index model (2), we use results from Gittins (1989) that explicitly calculate Gittins indices, G , for Bernoulli bandits with conjugate (Beta-distributed) priors. Specifically, given a Beta-distributed prior, a discount rate, α , and numbers of wins and losses prior to t , ω_{t-1}^i and ℓ_{t-1}^i , the index for an arm is a function of the triple $(\alpha, \omega_{t-1}^i, \ell_{t-1}^i)$. Table 11 in Gittins (1989) lists $G(\alpha, \omega_{t-1}^i, \ell_{t-1}^i)$ for arms in which $\alpha = 0.99$ and ω_{t-1}^i and ℓ_{t-1}^i range from 1 to 40.¹²

In calculating each arm’s Gittins index for any experimental trial, we assume the prior is Beta distributed and use the results from Gittins (1989). Given the prior information we communicate to subjects – that $\alpha = 0.99$ and that both arms had won twice and lost once in three prior trials – we further assume that $(\alpha = 0.99, \omega_0^i = 2, \ell_0^i = 1)$. After each choice of an arm and outcome, that arm’s ω_t or ℓ_t is updated and its G recalculated.

Note that the use of Beta-distributed priors is purely for computational tractability, and in the experimental sessions we did not inform subjects that the prior distribution was of this form. Thus, our experiment does *not* provide a sharp test of whether or not subjects are rational. Rather it tests how well the “rational model” with the given Beta prior fits subjects’ choices. This is in keeping with our original aim of validating models for use in the context of competitive analysis. To emphasize the distinction between a test for rationality and a test for model fit, we call the model “Gittins” (rather than “Rational”) when we report our experimental results.

Furthermore, even if we had informed subjects, in words or through pictures, that the prior is of a Beta ($\omega_0^i = 2, \ell_0^i = 1$) form, neither the prior distribution nor the discount rate *implicitly* used by the subject is observable. In theory, we might estimate ω_0^i , ℓ_0^i , and α to accommodate subjects’ unobserved, idiosyncratic priors and discount rates. The result would be a generalized model with three free parameters, and we would search over all feasible triplets $(\alpha, \omega_0^i, \ell_0^i)$ to find the initial parameters that generate the least inconsistency between model and experimental data. In practice, however, the calculation of the Gittins index is burdensome in and of itself, and we have not attempted to search among this broader class of Gittins-index policies.¹³

Myopic Model

¹²For ω_{t-1}^i or ℓ_{t-1}^i greater than 40, Gittins (1989) provides an approximation – equation (7.16), fitted with parameters from Tables 12–14 – that allows for calculation of indices that are typically precise within four decimal places.

¹³There exist closed-form approximations to the Gittins index that could be used to fit idiosyncratic subject prior and discount-rate information (Chang and Lai 1987, Brezzi and Lai 2002). The expressions are not accurate for discount rates that are significantly less than one, however.

For the Myopic model we, again, use Beta priors for convenience, and the resulting calculations are straightforward. Specifically, (3) does not require a discount rate, α , and given a Beta prior, the index is a straightforward function of previous wins and losses:

$$M(\omega_{t-1}^i, \ell_{t-1}^i) = \frac{\omega_{t-1}^i}{\omega_{t-1}^i + \ell_{t-1}^i}. \quad (9)$$

Given this form, it is not difficult to incorporate the effect of changes in the prior distribution on the indices.

Therefore, we test two versions of the myopic model. The first includes no free parameters and assumes $(\omega_0^i = 2, \ell_0^i = 1)$. We call this the “Myopic-0” model. The second incorporates one free parameter. Specifically, we define a common ω_0^i for both arms and let it range between 0.01 and 2.99, in increments of 0.01. We then define a common $\ell_0^i = 3.0 - \omega_0^i$. Thus, the second version still requires that both arms have the same initial prior, but it allows the shape of the prior to vary. We call this model “Myopic-1.”

Two elements of the parameter range are worth noting. First, to be consistent with the prior wins-loss information we report to subjects (2 wins in 3 prior trials), we bound ω_0^i away from 0 and 3, which respectively reflect beliefs that the probability of a win is zero and one. Second, by requiring $\omega_0^i + \ell_0^i = 3.0$, we fix the “strength” of the initial prior to be consistent with the quantity of prior information we report to subjects: the results of three prior trials.¹⁴

Simple Model

Recall that the Simple model hypothesizes that customers think of arms as being good or bad, with expected rewards of μ_G and μ_B . Algebraic manipulation (provided in Appendix E) demonstrates that we can write the index for arm i at time t as

$$\tilde{S}(\omega_{t-1}^i, \ell_{t-1}^i) = \omega_{t-1}^i \cdot \mathcal{U} - \ell_{t-1}^i \cdot \mathcal{D}, \quad (10)$$

the result of a series of ω_{t-1}^i “up-steps” ($\mathcal{U}$) and ℓ_{t-1}^i “down-steps” ($\mathcal{D}$). Furthermore, without loss of generality, we can normalize $\mathcal{U} \equiv 1$. The result is a model with one free parameter, $\mathcal{D}$.

For arm i with probability μ_i (not necessarily μ_G or μ_B) of winning, the expected “drift” of the random walk $\tilde{S}$ at time t equals $\mu_i \mathcal{U} - (1 - \mu_i) \mathcal{D}$. Then given $\mathcal{U} \equiv 1$ and some fixed $\mathcal{D}$, the subject’s “aspiration level” – the average quality level required of i so that the drift is non-negative – is

$$\mu^* = \frac{\mathcal{D}}{\mathcal{D} + 1},$$

and once a Simple subject begins sampling from an arm with $\mu_i \geq \mu^*$, her expected switching time is infinite. (See the discussion in Section 3.2, as well as Gans (2002b).) When estimating the fit of

¹⁴This setup implicitly assumes that, before being informed of two successes in three prior trials, subjects have noninformative priors. An alternative, which we have not tested, would be to let $\omega_0^i = \gamma + 2$ and $\ell_0^i = \gamma + 1$ for both arms, where $\gamma > -1$ is the model’s single free parameter.

subjects' choices with the Simple model, we systematically vary $\mathcal{D}$ so that μ^* ranges from 0.0033 to 0.99 in even increments of 0.0033.

We bound $\mathcal{D}$ away from zero so that the model is required to penalize arms for bad outcomes. In contrast, for $\mathcal{D} = 0$, the Simple model recommends an the arm with the greater number of wins – without regard to numbers of losses – and would provide (perhaps unfairly) good fits for subjects that never change arms, no matter how many losses. (For example, see the results for HH- n in §6.3.2, below.)

Last n

As with the previous models, for convenience we use Beta prior, so at trial t the index for arm i is the ratio of the number of wins and losses in the previous n trials on the arm. Formally, we let $\omega_t^i(n)$ and $\ell_t^i(n)$ be the number of wins and losses in the last n trials on arm i , so that $\omega_t^i(n) + \ell_t^i(n) = n$, and

$$L(\omega_{t-1}^i(n), \ell_{t-1}^i(n)) = \frac{\omega_{t-1}^i(n)}{\omega_{t-1}^i(n) + \ell_{t-1}^i(n)}. \quad (11)$$

While we have tested models for $n = 1, 2, 3, 4, 5$, we report results only for $n = 1, 3, 5$. This allows the figures to be less cluttered and easier to read, and the omitted results (for $n = 2, 4$) are consistent with the broader trends seen across $n = 1, 3, 5$.

We have also tested a “meta” model which treats n as an additional free parameter and uses the n associated with each subject's lowest BIC score. The results of this meta-model are not fundamentally better than those of the basic Last- n family.

Finally, we note that, at $t = 0$ the Last- n model requires data for periods $t = -1, -2, \dots, -n$. For simplicity, we have initialized the record of these prior outcomes to all be wins. While this assumption is not consistent with the prior information shown to the subjects, it is consistent with the initial conditions required for fitting the Hot Hand family of models. (See below.) As we will see in the next subsection, differences in fit are substantial across models, and we do not believe that these initial conditions have significantly affected our results.

Hot Hand

Recall that the Hot Hand family of models uses only the results of the most recently sampled arm to decide which arm to sample next. Specifically, if the n previous trials on the current arm were all losses, then the HH- n model recommends switching to the other arm; otherwise the model recommends continuing to sample from the current arm.

While the Hot Hand rule is *not* index-based, when fitting the model to subjects' observed choices, it will be convenient for us to define it as an index rule. Therefore, we formally define HH- n indices as follows. For a subject that sampled from arm i at time $(t - 1)$, we define the indices for arms i and

$j \neq i$ to be

$$HH^i(\ell_{t-1}^i(n)) = 1\{\ell_{t-1}^i(n) < n\} \quad \text{and} \quad HH^j(\ell_{t-1}^j(n)) = 1 - HH^i(\ell_{t-1}^i(n)), \quad (12)$$

where $\ell_{t-1}^i(n)$ denotes the number of losses in the previous n contiguous trials on arm i . The rule then recommends choosing the arm with the larger of the two indices.

If at time t a subject chooses an arm, i , that she had chosen in the previous trial, then $\ell_t^i(n) = 1\{i \text{ loses at } t\} + \ell_{t-1}^i(n)$. Because the counter $\ell_{t-1}^i(n)$ only tracks the number of losses in the previous n contiguous trials on arm i , we reset $\ell_{t-1}^i(n) = \ell_{t-1}^j(n) = 0$ whenever the subject switches arms: that is, whenever $HH^i(\ell_{t-1}^i(n)) \neq HH^i(\ell_t^i(n))$. This implies that HH- n recommends staying on the current arm whenever the current run on an arm is less than n trials.

Because the prior information we provide to subjects at the start of each session does not distinguish the order in which “prior” samples of the two arms were made, it is not well determined whether a subject’s first trial in a given session represents a switch to a new arm or the continuation of a run on the current arm. Therefore, for simplicity, we assume that the first trial represents a switch, and we reset the associated loss counters, $\ell_0^i(n) = \ell_0^j(n) = 0$, accordingly.

Exponential Smoothing

The Exponential Smoothing model (7) is implemented in a straightforward fashion. We test two versions of it.

The first model (ES-1) has the smoothing weight, γ , as its one free parameter. The initial index of each arm is fixed at $ES_0^i = 2/3$, so that it matches the win-loss ratio reported as prior information. Note that, for $\gamma = 1$, ES-1 corresponds to a Last-1 model, so to better distinguish between the two models, we bound γ away from 1. Similarly, for $\gamma = 0$, both arm’s indices would equal $2/3$ for all t and would not be informative. Therefore, we vary γ from 0.01 to 0.99 in increments of 0.01.

The second model (ES-2) also treats the initial index ES_0^i as a free parameter. In this case we vary both γ and ES_0^i from 0.01 to 0.99 in increments of 0.01, for a total of roughly 10,000 possible combinations of free parameter values.

6.2 Fitting the Models to Subjects’ Choices

Let the generic index, I , represent the index of the model being used, and suppose that, at trial t in a given session, a model’s indices are I_t^l and I_t^r for the left and right arms. If $I_t^l > I_t^r$, then a subject whose choices are dictated by the model will choose the left arm. In fact, at t , the subject will either choose the left arm or not, and a straightforward and readily observable measure of consistency would simply record whether or not the subject’s choice matches the prediction or not.

An aggregate measure of consistency over all 347 trials would then be the total number of incorrect model predictions, the smaller the number the better. The determination of this number is trivial for

the Gittins, Myopic-0, Last-n, and HH-n models, since they have no free parameters. For the Myopic-1, Simple, ES-1 and ES-2 models, one can search for the values of the respective free parameters that minimize the total number of errors across all 347 trials.

Alternatively, at any given trial one might further judge *how well* a subject’s choice matches that prescribed by a model, rather than simply whether or not the observed choice is consistent. A common means of judging the degree of consistency is through the use of random utility models (Anderson et al. 1992). Here, one posits that the value of a left or right choice at a particular trial is randomly distributed and that only the mean of the distribution is captured by the model indices I_t^l and I_t^r . In the context of our experiments, the random fluctuation might be ascribed to errors (or “trembles”) in judging the expected value of the choice.

If, from trial to trial, this random noise is independently and identically distributed according to a Gumble (double exponential) distribution then we have a so-called logit model. In this case, the probability of choosing the left arm at trial t is

$$P\{\text{choose left}\} = \frac{e^{\beta I_t^l}}{e^{\beta I_t^l} + e^{\beta I_t^r}}. \quad (13)$$

Similarly, $P\{\text{choose right}\} = 1 - P\{\text{choose left}\}$.

Note that, by nesting the original choice models within the logit framework, we have imposed an additional (and unobservable) level of complexity, as well as the addition of a free parameter, β . The benefit is that we have a means of judging how well each observed choice matches a model’s prediction. Furthermore, that measure is a probability, and we can easily add the 347 choices’ log-probabilities to calculate an aggregate log-likelihood (LL) of observing each subject’s outcome.¹⁵

A likelihood measure of consistency is particularly appealing in that it also lets us naturally correct for differences in the numbers of free parameters used by the models, something which is not easily accomplished when counting the total numbers of errors. We account for free parameters using the Bayesian Information Criterion (BIC) $\text{BIC} = -2\text{LL} + \text{d.f.} \times \ln(347)$, where d.f. is the numbers of degrees of freedom used by the model’s free parameters. Indeed, BIC scores have the appealing property that they can be used to approximate Bayes factors, the posterior odds the data being explained by one non-nested model, rather than another (Kass and Raftery 1995).

In contrast, a straightforward count of the number of inconsistent choices can become problematic if the index takes on only a few values. For example, in the extreme case that a choice model has $I_t^l = I_t^r$, regardless of a subject’s observed choices and outcomes, the number of inconsistent choices is always equal to zero, and the model “perfectly” fits the observed data. In contrast, the likelihood approach would record that each choice is completely random ($P\{\text{choose left}\} = P\{\text{choose right}\} = 0.5$,

¹⁵Our implementation of the Hot Hand model defines the index of the recommended arm as one and that of the other arm as zero. While this choice appears to be somewhat arbitrary, we note that it is without loss of generality, since the β of the superimposed logit model acts as an independent scaling factor. Therefore, the LL derived from the Hot Hand model is independent of the scale of its indices. (Thanks to Ed Kaplan for pointing this out.)

without regard to β), and the associated BIC score would, more appropriately, suffer. For this reason, we emphasize results obtained from using the logit model.¹⁶

The calculation of each model’s LL requires the solution of a nonlinear optimization problem. For models with no free parameter, the LL is concave in β and the optimal β can be found using standard optimization software. While the LL’s of the Myopic-1, Simple, and ES models are also concave in β , they are not necessarily jointly concave in β and their free parameters. Therefore, for these models, we perform a nested optimization procedure: the top level is a grid search over each model’s original free parameters; then, for each of these grid values, we use a solver to find the optimal β and LL. We then record the best LL among all of the top-level grid values.

We implement this procedure for each subject to find one best set of free parameter values across all 347 trials. For each subject, we use the same free-parameter values across all three sessions. This approach reflects our assumption that each subject’s parameters should be stable across treatment conditions, and it reduces the possibility of model overfitting.

Given the large quantities of data that we analyze, the use of BIC scores is also computationally appealing. In total, we have calculated more than 2 million scores for model fit – more than 10,500 free-parameter values associated with the various models for each of 195 subjects – and this task has been facilitated by the straightforward optimization required by log-likelihood and BIC scores.

6.3 Results

For each of 195 subjects, we have used the optimization procedure described above to find free parameters that minimize each model’s BIC score for each subject, and we report the fit results below. Because of the large number of models considered, we report the results of the various models by family: first those that are derived from the Gittins index, then the Last- n group, then Hot Hand models, and finally Exponential Smoothing.

6.3.1 Models Derived from the Gittins Index

Figure 3 presents two pictures of the BIC results for the Gittins-index family of models. The left panel displays an aggregate, across-subject view of the models’ performance, and the right panel shows the results of within-subject rankings of BIC scores.

More specifically, the curves in the figure’s left panel represent cumulative distributions of BIC scores.¹⁷ From the plot we see that Simple and Myopic-1, the models with an extra free parameter,

¹⁶We report results in the context of the models’ BIC scores. For the Gittins-index family of models, we have also performed an analogous analysis using numbers of errors as the measure of fit, and the results are largely the same.

¹⁷To derive a given model’s curve we sort its 195 BIC scores from lowest to highest – that is, best to worst. Then we plot the number of subjects with BIC scores less than or equal to each value listed on the x -axis. Models whose curves are farther “up and to the left” are interpreted as better fitting the observed data.

Figure 3: BIC Scores for Models Derived from the Gittins-Index

nearly dominate Myopic-0 and Gittins, even after the BIC score penalizes these models for the added parameter. Careful inspection shows that the Simple model does not, in fact, dominate the rest of the models in the 310-to-340 and the above-440 range, however. In the above-480 range, where the absolute fit of any of the four models is very poor, Myopic-0 outperforms Simple and Myopic-1, since its BIC score is not penalized for an extra free parameter.

The total BIC scores shown in the left panel's inset are consistent with these curves. When individual BIC scores are summed across individuals, we find that, in aggregate, the Simple model has the best performance, then Myopic-1, then Myopic-0, and finally the Gittins index itself.

We can also judge the quality of fit at the individual level. For every subject, we rank each model's fit, by BIC score. For each model, the first shaded bar is the number of subjects for whom it was the first-best fit. The next white part is the number of subjects for whom it was the second-best fit, and so on. The right panel of Figure 3 displays these results.¹⁸

The results show that, on a subject-by-subject basis, the Simple model fits the observed choices best most often, in about 41% of all subjects. Again, the Simple model has roughly twice the number of first place fits when compared to Myopic-0 or Myopic-1. In this case, however, the Myopic-0 model slightly outperforms Myopic-1.

We also performed pairwise comparisons of the models across subjects using paired- t and Wilcoxon Signed Rank tests. In all but one case, the differences were significant, with desired p -values vanishingly small. For the one-sided (alternative) hypothesis that the Myopic-1 BIC > Simple BIC, the results were somewhat less strong, however. Here, the t -test resulted in p -value of 0.08 while the signed rank test a p -value of 0.11.¹⁹

¹⁸We note that, in the figure, a tie for equal BIC scores results in all relevant policies receiving the higher rank. Thus, while each policy's rankings add up to $n = 195$, adding up a given ranking (e.g. 1st place) across models may lead to a total that exceeds or falls short of 195.

¹⁹Of the 195 subjects included in the test, there was one extreme outlier, Subject 127. After removing this outlier

6.3.2 Other Choice Models

We next present analogous results for the Last- n , hot Hand, and Exponential Smoothing models. To facilitate comparison with the Gittins-index family of models described above, the figures include results for the Simple model as well.

Last- n

The left panel of Figure 4 shows that the BIC scores associated with the Last- n family of models tend to be higher than those associated with the Simple model. Both the cumulative distributions and the aggregate BIC scores of the Last-5 and Last-3 models are dominated by the Simple model's. The Last-1 model outperforms the Simple model only when BIC scores are high and none of the models fits particularly well.

Figure 4: BIC Scores for Last- n Model

In contrast, the figure's right panel shows that, on a subject by subject basis, Last-1 appears to perform nearly as well as the Simple model. While a one-sided paired t -test, with an (alternative) hypothesis that Last-1 BIC > Simple BIC, was significant at nearly the 0.01 level, an analogous Wilcoxon Signed Rank test returned a weaker p -value, of 0.111.

As the left panel suggests, however, subjects with lower Last-1 BIC scores are those for whom neither model fits very well. For example, for the 88 subjects for whom the Simple ranked first, the average Simple model's BIC score was 328.7, while for the 79 subjects for whom the Last-1 model ranked first, the analogous score was 392.4. In fact, in cumulative distribution of the 1st-ranking BIC scores, the Simple model also dominates Last-1.

Hot Hand

Figure 5 details the BIC scores of the Hot Hand family of models. One sees that, in contrast to the from the data, the p -value for the t -test improved to 0.03, and that for the signed rank test to 0.08.

Last- n model, the Hot Hand models' fit to subjects' choices improves with larger n .

HH-5, in particular, appears to perform extremely well on both an aggregate and an individual basis. The left panel shows that the model's aggregate BIC score dominates that of the other models, as does the cumulative distribution of its BIC scores. The right panel shows that, on a subject-by-subject basis, HH-5 significantly outperforms the Simple model as well. Both a one-sided paired t -test and a Wilcoxon Signed Rank test that compare HH-5 and Simple BIC scores returned vanishingly small p -values.

Figure 5: BIC Scores for the Hot Hand Model

Note also that, within the Hot Hand family, the performance of the models is well-ordered. In both of Figure 5's panels, HH-5 outperforms HH-3, and HH-3 outperforms HH-1. Indeed, while we do not report detailed results here, the results for HH-4 and HH-2 also fit within this ordering.

To better understand why the rankings appear to be ordered with n , we compared the models' BIC scores to subjects average run lengths. The results, shown in Figure 6, are informative: the HH-5 model's BIC scores are strongly associated with average run length, with a nearly linear relationship between BIC and $\bar{\tau}$ for $\bar{\tau} \leq 10$.²⁰ Close inspection shows that a more attenuated form of the same type of relationship also holds for HH-3, and plots that include the results for the HH-2 and HH-4 models show a consistent progression: the larger the n , the more strongly average run length is associated with low BIC scores. This leads us to question the quality of the Hot Hand model's fit to the data.

More generally, consider a sample path in which an arm with a probability of p of winning is pulled over and over again. Then it can be shown that the expected fraction of trials in which a HH- n model recommends switching is of order $O((1-p)^n)$. (See Appendix C.) As $n \rightarrow \infty$, the frequency of recommended switches approaches zero exponentially quickly.

²⁰To quantify the strength of the relationship, we also created a binary variable that took on a value of 1 for a subject if HH-5 was ranked first and zero otherwise. Tests for the probability of HH-5 obtaining a 1st ranking showed that $\bar{\tau}$ has, in fact, a significant positive impact. In a logit regression with an intercept, the coefficient for $\bar{\tau}$ was positive with a t -statistic greater than 4.6. A similar probit regression returned the same direction and significance.

Figure 6: Relationship Between BIC Scores and $\bar{\tau}$ in the Hot Hand Model

Now suppose that in T trials a subject switches arms m times. If an HH- n with very large n is fit to the data, it will register roughly m inconsistencies, perhaps fewer. In turn, if m is small then the fit associated with the HH- n model will be very good, no matter what the relative winning probabilities of the two arms. Therefore, in a significant sense, the HH- n model may be obtaining good fits simply by picking up long run lengths, and we suspect that the BIC scores of HH-6, HH-7, etc. would keep improving.

As a simple test of this hypothesis, we tested a HH- ∞ model that recommends *never* switching. Its aggregate BIC score was 61,659, much better than that of HH-5. Furthermore, as the plot of BIC on $\bar{\tau}$ in Figure 6 confirms, HH- ∞ displays a more extreme version of the pattern seen in the other Hot Hand models. Thus, we believe that the low BIC scores obtained by HH-5 are largely artifacts of subjects' long run lengths.

In contrast, there is some evidence against this phenomenon holding for the other models. Analogous plots for the other models do not show the same strong relationship between BIC scores and $\bar{\tau}$. (See Appendix F.) Rather, BIC scores tend to be higher for small average run lengths and then drop for average run lengths above $\bar{\tau} \approx 7$. Above this cutoff, there appears to be only a mild, negative relationship between the two. An exception to this general statement is the Last-1 model, whose BIC scores are best for subjects with very low $\bar{\tau}$ s. Again, these are the subjects who also have high BIC scores and whose switching behavior seems to be nearly independent of prior outcomes.

Exponential Smoothing

Lastly, Figure 7 compares the performance of Exponential Smoothing to that of the Simple and Myopic-1 models.

Figure 7: BIC Scores for Exponential Smoothing Model

The figure’s left panel shows that, in the aggregate and for most of the range of the BIC scores, the ES-2 model outperforms the ES-1, the Simple, and the Myopic-1 models. In the upper range of BIC scores – the cases in which model fit is poor – the ES-2 model is tied with the ES-1 model. While the ES-1 model outperforms the Simple and Myopic-1 models, both in aggregate and in the above-325 range of BIC scores, all three models are dominated by ES-2.

The rankings in the right panel show that, on a subject-by-subject basis, the BIC scores of ES-2 dominate those of the other models. Similarly a one-sided paired t test comparing ES-2 and ES-1 scores yielded a vanishingly small p-value, and an analogous Wilcoxon Signed Rank test returned a p-value of 0.01. Tests among other pairs of models – ES-2 versus Simple, ES-1 versus Simple, and ES-1 versus Myopic-1 returned p-values that ranged from highly significant (less than 0.001) to vanishingly small.

6.3.3 Summary of Results

Recall that we are interested in two dimensions of model performance. The first is fit to subjects’ observed choices, as measured by both aggregate and individual-level comparisons of models’ BIC scores. The second is tractability, the ability to derive simple (preferably closed-form) expressions for aggregate statistics – such as or expected switching time or fraction of times chosen – from the models’ primitive parameters. Table 1 summarizes our findings.

Among the Gittins-index family of models, the Simple model performed best. While the BIC scores of Myopic-1 were on par with those of the Simple model, the Simple model is more tractable. The original Gittins-index model of “rational” choice performed most poorly, both in terms of both

	Gittins-Index Family				Descriptive Models			
	G	M-0	M-1	S	Last-1	HH-5	ES-1	ES-2
Aggregate BIC	83,236	79,344	77,304	76,776	80,278	66,910	73,493	69,018
Tractability	--	-	-	+	++	++	+	+

Table 1: Summary of Model’s Fit to the Data and Tractability

fit and tractability.

Among the descriptive models, the Hot-Hand 5 performed best. Indeed, Hot Hand dominated both the Gittins-index family and the other descriptive model in terms of fit, and it is also easy to analyze in terms of deriving useful aggregate statistics. The strong performance of the Hot Hand model must be qualified, however. As Figure 6 shows, the model’s low BIC scores appear to be artifacts of longer run lengths.

There are other limitations to the Hot Hand model as well. While the Hot Hand’s summary performance measures are easy to analyze in the context of Bernoulli outcomes, the model does not immediately apply to more complex reward distributions.²¹ In contrast, the other models we have tested can be applied directly to generally distributed rewards. More fundamentally, HH- ∞ is not applicable to analysis of supplier competition, since it models consumers as not responding to the quality of service they receive.

Of the second set of models that we tested, Exponential Smoothing also performed quite well, both in terms of analytical tractability and in terms of model fit. Exponential Smoothing compares well with the Simple model, with (roughly) the same level of analytical tractability and stronger BIC scores. We therefor favor the use of the Exponential Smoothing model, and we elaborate on this recommendation, below.

7 Discussion

Our experimental results are positive on two levels. First, they demonstrate that the expected number of trials on a given arm is increasing and convex in the arms’ expected reward. This finding is consistent with the models of choice we have considered, as well as with more general industry observation (e.g., Jones and Sasser (1995)). Second, they support the use of more analytically tractable models for use in the type of competitive analysis pursued in Gans (2002a), Gaur and Park (2003), and Hall and Porteus (2000). Both the Exponential Smoothing and Simple models, in particular, are analytically tractable, flexible enough to be used with generally-distributed rewards, and robust with respect to model fit.

²¹One can generalize the Hot Hand model to handle rewards that are not Bernoulli distributed by considering “wins” to be outcomes above some threshold and “losses” to be outcomes below. This threshold would then become a free parameter to be searched for.

We note that not all of the less elaborate models performed better in measures of fit. For example, the Last- n model fit the data poorly for higher n 's, and the HH- n model performed poorly for lower n 's. Thus, simplicity, by itself, does not appear to be a guarantee of a model's success.

The Exponential Smoothing model particularly appeals to us for a number of reasons. First, it appears to fit the data well. Furthermore, the special case of $\gamma = 1$ – which was excluded from consideration in our fit analysis – corresponds to a Last-1 model, so the model is (marginally) even more flexible than the results already indicate. Also, recall from Section 3.3 that Exponential Smoothing is consistent with a model of a Bayesian subject that faces Markov-modulated rewards. We find this representation to be intuitively satisfying. Even though operations management papers often model the world as being stationary it is not; the ES models implicitly capture a subject's belief that her reward system is always changing.

Still, the fact that the HH- ∞ model achieves the lowest BIC scores is, at first glance, discomfiting, and one may wonder if there is something about our experimental setup or analysis that should be changed to more properly penalize this type of “artifactual” performance.²² We believe it will be difficult to eliminate this effect. Because the BIC score is an adjusted measure of the likelihood of observing a given sample path, samples with low base rates of switching tend to provide good fits for HH- ∞ . Conversely, HH- ∞ should have a poor fit for subjects who switch frequently. Therefore, to properly penalize the HH- ∞ model's behavior, one needs to set experimental conditions so that the observed switching rate among arms is higher.

How does one generate higher switching rates? The descriptive statistics reported in Section 4 suggest that more frequent switching is associated with lower win rates and small, rather than large, differences between the two arms' $P\{\text{win}\}$ s. In addition, we hypothesize that a smaller expected number of trials would increase switching, since a greater fraction of a subject's total winnings would depend on each trial.

Shorter sessions with high overall rates of switching would provide fewer, less discriminatory, trials over which the models' differences can be highlighted, however. That is, an experimental design with lower win rates and fewer trials is likely to push all models' BIC scores up towards that of random switching and make it more difficult for the various models' likelihoods to diverge. Therefore, we believe that a design including many trials with intermediate-level win rates that we used is a reasonable, if not perfect, experimental choice.

The use of a large number of trials has also allowed us to more clearly distinguish among the performance of models with various levels of learning. With many trials, models with more complete

²²One may also wonder whether the Simple and Exponential Smoothing models' free parameters reflect artifacts that are similar to the Hot Hand's. We have plotted both BIC score and average run length against free parameter values, and the results do not appear to reflect relationships that are more substantial than the original results for BIC scores and $\bar{\tau}$. For plots of BIC scores against $\bar{\tau}$, see Appendix F, and for plots of BIC scores against free parameter value, see Appendix G.

learning, such as the Gittins and Myopic models, tend to converge on an arm. Conversely, the Exponential Smoothing, Last- n , and Hot Hand models represent a much weaker forms of learning, never converging on an arm. The Simple model represents something of a median, converging only in expectation, and only then if an arm meets the aspiration level μ^* .

The BIC scores suggest that subjects' choices are more consistent with models of choice that are both more myopic and attenuated in their form of learning. Aggregate measures of learning are also consistent with this finding. For example, on average, subjects switched 11.9 times in the 1st half of each session and 9.2 times in the 2nd half. While this 22.4% decrease is important, it is far from reflecting a general convergence to a single arm.

While the Gittins index fared poorest in terms of BIC score, we recall that computational limitations have forced us to fix subjects' implicit (unobserved) discount rates to match the $\alpha = 0.99$ we used in the experiments. One may ask how much better the Gittins index would have performed had we had the ability search for the best-fitting implicit discount rate for each subject. A partial answer can be found by recalling that Myopic-0 is simply a Gittins index model with $\alpha = 0$. The fact that Myopic-0's BIC scores dominates the Gittins index's suggest that subjects *are* significantly more myopic than is optimal. Still, the questions concerning how much better the Gittins index would perform with an intermediate α remains open.

Section 6's results also suggest that there may be multiple segments of subjects, some of whose choices are more strongly guided by one model or another. This is an important issue for future experimental analysis, as well as for the application of our results. The fact that various segments of customers base their response on different choice models affects overall market outcomes. Analogously, for representations with free parameters, such as the Simple and Exponential Smoothing models, the distribution of parameter values determines overall market response. Thus, the market response functions which are derived from these choice models – and used in operations management research – should, ideally, account for these differences.

Finally, we note that, in addition to their trial-by-trial choices, subjects provided data in brief end-of-session and end-of-experiment questionnaires. While we have not yet sought to match subjects' best-fitting models with other covariates, there may exist data that make this analysis possible.²³

²³We have looked at demographic covariates, which do not appear to be associated with subject performance. An OLS regression of earnings on the experiment's duration and $\bar{\tau}$ – controlling for demographics provided such as school, age, gender, ethnicity, and whether they reported they were confused by the experiment – shows the same results as a regression in which when we do not control for demographics: a unit increase in $\bar{\tau}$ is associated with an earnings increase of \$0.07; and a one-minute increase in the experiment's duration is associated with an earnings' increase of \$0.03.

8 Conclusion

In this paper, we have examined the descriptive ability of simple choice models in a bandit setting. We first developed a hierarchy of Bayesian models that range widely in complexity, from the Gittins-index model that most closely represents rational behavior, though the simpler Myopic model, to the Simple model in which the customer’s view of a supplier is representative of a category. We also considered other well-known choice models: the Last- n models, which are limited-memory analogue of the Myopic model; the Hot Hand models which respond to sequences of losses on the current arm; and Exponential smoothing models, which take simple weighted averages of current and prior information.

We performed two sets of analysis, each intended to test the correspondence between subjects’ behavior and the models’ predictions. The first confirmed that the average run lengths on an arm appears to be increasing and convex in the arm’s average reward. The second demonstrated that subjects’ actual choices more closely matched the recommendations of less complex representations of choice.

Our results may also be of use in the larger task of building a positive theory of dynamic decision making, as described in Hutchinson and Meyer (1994). They suggest that there may be segments of subjects, some of whose choices are more strongly guided by one model or another. Similarly, a given subject may use more than one model or choice, perhaps shifting strategies over time. While we have not approached subjects’ behavior in this fashion, the data exist that may make the analysis possible. Additional models, such as Q-learning, can be tested, as can more exploratory approaches. A recent paper by Houser et al. (2004) provides an exciting, new method for more exploratory analysis.

At the same time, there exist a number of limitations to our findings that bear repeating. First, while the experimental environment was essential for controlling the attributes of the arms and for generating predictions, as in any experiment, the artificiality of setting may lead to behavior that differs from the way people solve real world bandit-like problems. More narrowly, different experimental parameters, such as the expected number of trials $((1 - \alpha)^{-1})$, probabilities of success, or reward distributions may allow us to further discriminate among models’ performance. More generally, the bandit model, itself, may mis-specify the basis of customer choice, though this (of course) was not a modelling problem these experiments were intended to address. For example, it may not be reasonable to assume that the quality distributions of the arms set by suppliers are independent.

Nevertheless, on the whole these results are positive for researchers in operations management. We originally viewed the models being tested as representing different tradeoffs between analytical tractability and richness of the representation of the learning process. The experimental results suggest that subjects’ learning is less strong than might be expected, however. It appears that the more tractable models that have been favored in operations management research can more than “adequately” capture the essentials of customer choice behavior.

References

- C. M. Anderson. *Behavioral Models of Strategies in Multi-Armed Bandit Problems*. Ph.D. Dissertation, California Institute of Technology, 2001.
- S. P. Anderson, A. de Palma, and J-F. Thisse. *Discrete Choice Theory of Product Differentiation*. MIT Press, 1992.
- J. Banks, M. Olson, and D. Porter. An Experimental Analysis of the Bandit Problem. *Economic Theory*, 10:55-77, 1997.
- M. Brezzi, and T. L. Lai. Optimal Learning and Experimentation in Bandit Experiments. *Journal of Economic Dynamics & Control*, 27:87-108, 2002.
- R. G. Brown and R. F. Meyer. The Fundamental Thoerem of Exponential Smoothing. *Operations Research*, 9:673–687, 1961.
- B. D. Burns and B. Corpus. Randomness and Inductions from Streaks: “Gamblers Fallacy” Versus “Hot Hand.” *Psychonomic Bulletin & Review*, 11:179–184, 2004.
- R. R. Bush and F. Mosteller. *Stochastic Models for Learning*. Wiley, New York, 1955.
- C. F. Camerer and T-H. Ho. Experience-Weighted Attraction Learning in Normal Form Games. *Econometrica* 67:827–874, 1999.
- C. F. Camerer and R. M. Hogarth. The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. *Journal of Risk and Uncertainty*, 19:7–42, 1999.
- F. Chang and T. L. Lai. Optimal Stopping and Dynamic Allocation. *Advances in Applied Probability*, 19:829-853, 1987.
- T. Erdem and M. P. Keane. Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets. *Marketing Science*, 15:1-20, 1996.
- I. Erev and A. Roth. Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed-Strategy Equilibria. *American Economic Review*, 88:848–879, 1998.
- D. Friedman and S. Sunder. *Experimental Methods: A Primer for Economists*. Cambridge University Press, 1994.
- N. Gans. Customer Loyalty and Supplier Quality Competition. *Management Science*, 48:207-221, 2002a.
- N. Gans. Service Quality, Switching Costs and Their Effect on Customer Retention. OPIM Department Working Paper, The Wharton School, University of Pennsylvania, 2002b.

- V. Gaur and Y-H. Park. Asymmetric Consumer Learning and Inventory Competition. Working Paper, NYU Stern School of Management, 2003.
- I. Gilboa and A. Pazgal. Cumulative Discrete Choice. *Marketing Letters*, 12:119-130, 2001.
- T. Gilovich, R. Vallone, and A. Tversky. The Hot Hand in Basketball: On the Misperception of Random Sequences. *Cognitive Psychology*, 17:295–314, 1985.
- J. C. Gittins and D. M. Jones. A Dynamic Allocation Index for the Sequential Design of Experiments. *Progress in Statistics*, J. Gani et al. eds., 241-266, 1974.
- J. C. Gittins. Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society*, B41:148-177, 1979.
- J. C. Gittins. *Multi-Armed Bandit Allocation Indices*. John Wiley & Sons, 1989.
- P. M. Guadagni and J. D. C. Little. A Logit Model of Brand Choice Calibrated on Scanner Data. *Marketing Science*, 2:203-238, 1983.
- J. Hall and E. Porteus. Customer Service Competition in Capacitated Systems. *Manufacturing & Service Operations Management*, 2:144–165, 2000.
- D. W. Harless and C. F. Camerer. The Predictive Utility of Generalized Expected Utility Theories. *Econometrica* 62(6):1251-1289, 1994.
- G. W. Harrison. Theory and Misbehavior of First-Price Auctions. *American Economic Review*, 79:749–762, 1989.
- P. W. Henderson and R. A. Peterson. Mental Accounting and Categorization. *Organizational Behavior and Human Decision Processes*, 51:92–117, 1992.
- A. D. Horowitz. *Experimental Study of the Two-Armed Bandit Problem*. Ph.D. Dissertation, University of North Carolina, Chapel Hill, 1973.
- D. Houser, M. Keane, K. McCabe. Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm *Econometrica*, 72:781–822, 2004.
- J. W. Hutchinson and R. J. Meyer. Dynamic Decision Making: Optimal Policies and Actual Behavior in Sequential Choice Problems. *Marketing Letters*, 5:369–382, 1994.
- T. Jones and E. W. Sasser, Jr. Why Satisfied Customers Defect. *Harvard Business Review*, 73(6):88–99, 1995.
- D. Kahneman and A. Tversky. On the Psychology of Prediction. *Psychological Review*, 80:237–251, 1973.

- R. E. Kass and A. E. Raftery. Bayes Factors. *Journal of the American Statistical Association*, 90:773-795, 1995.
- E. L. Lehmann. *Nonparametrics: Statistical Methods Based on Ranks*. Holden-Day, 1975.
- J. G. March. Learning to Be Risk Averse. *Psychological Review*, 103:309-319, 1996.
- T. Matsuda and M. Sekiguchi. Models of Human Forecasting Behavior: A Note on the Relationship Between Exponential Smoothing and the Bayesian Method. *Journal of the Operational Research Society of Japan*, 13:136-154, 1971.
- R. J. Meyer and Y. Shi. Sequential Choice Under Ambiguity: Intuitive Solutions to the Armed-Bandit Problem. *Management Science*, 41:817-834, 1995.
- G. A. Miller. The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity to Process Information. *Psychological Review*, 63:81-89, 1956.
- D. Robbins and P. L. Warner. Individual Organism Probability Matching with Rats in a Two-Choice Task. *Bulletin of the Psychonomic Society*, 2:405-407, 1973.
- A. E. Roth. Laboratory Experimentation in Economics: A Methodological Overview. *The Economic Journal*, 98:974-1031, 1988.
- R. Schmalensee. Alternative Models of Bandit Selection. *Journal of Economic Theory*, 10:333-342, 1975.
- H. A. Simon. Theories of Decision-Making in Economics and Behavioral Science. *American Economic Review*, 49:253-283, 1959.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- A. Tversky and D. Kahneman. Judgment under Uncertainty: Heuristics and Biases. *Science*, 185:1124-1131, 1974.

Appendix for

Simple Models of Discrete Choice and Their Performance in Bandit Experiments

Noah Gans	George Knox	Rachel Croson
OPIM Department	Marketing Department	OPIM Department

The Wharton School, University of Pennsylvania
Philadelphia, PA 19104

`{gans,gknox,crosonr}@wharton.upenn.edu`

Appendix

A Background on the Multi-Armed Bandit

A “one-armed bandit” is a slot machine that is found in casinos. To play it, a gambler puts a coin in the slot and pulls the machine’s arm (its handle), thereby spinning a set of flywheels. When the flywheels stop turning, the pictures that they display - for example, three cherries or two lemons and a bell - determine the gambler’s payoff. They also determine the starting position of the flywheels for the next pull of the arm.

A multi-armed bandit problem is a mathematical abstraction of the one-armed bandit. One may think of it as a set of one-armed bandits that is repeatedly played by a single gambler. At each gamble, she chooses one of the arms. A pull of an arm advances its state and yields a payoff, but it does not change the states of the other arms. If the distribution of states obtained by a pull of an arm depends only on the arm’s previous state, so that its evolution is Markov, then a multi-armed bandit with m arms can be modelled as a set of m Markov chains, each independent of the next. Associated with each state of an arm’s Markov chain is a payoff distribution.

A common objective for players of the multi-armed bandit is to maximize the expected discounted payoffs over some horizon, either finite or infinite. At any trial, the gambler compares the current states of the various arms and chooses the one that maximizes the expected value of her discounted future rewards. Note, however, that the arm that maximizes the expected discounted future payouts need not be the arm with the highest expected payout on the next pull. Indeed, the “best” arm to pull depends on the overall structure of each of the various Markov chains, as well as its current state.

The bandit framework can be used to model learning phenomena. The so-called Bayesian bandit assumes that each arm has a fixed but unknown payoff distribution and that the decision-maker has a prior (distributional) belief concerning which payout distribution each arm actually has. After choosing an arm and receiving the sample payout, the decision-maker revises her beliefs about the payout distribution using Bayes’s rule. In this case, the state of each arm is its prior distribution, and the choice of one arm does not affect the prior distribution of any other’s. Each payout is a bit of information that the decision-maker can use to better learn the true payout distribution of the arm that was pulled.

B Experiment Screen Shots

This appendix details the sequence of screens that a subject sees as she takes part in the experiment.

The first page explains the procedure of participating in the experiment over the internet, the maximum amount of time subject may take between decisions (15 minutes), and the contact information in case anything goes wrong.

[Webmaster email link](#)

Welcome!

In this exercise you will be making a series of choices, which will be presented to you in a sequence of web pages. Once you have made a given choice you cannot go back and change it. Thus, while you are completing the exercise, the 'back' button on your browser should not be used. Pressing that button will terminate the program. Once the program has terminated, you cannot rejoin or begin the exercise again.

If you proceed directly through the exercise with a fast internet connection, it should take you about 1 hour (slower connections will take longer). If you wish to take a short break between choices, you may, but the program is set to terminate in response to a delay of greater than 15 minutes. Once the program has terminated, you cannot rejoin or begin the exercise again.

You will be paid for your participation in cash. The amount of money you earn will depend on the choices you make, and will be paid to you after the exercise is completed, along with a fixed, \$5 participation payment. You will need to bring your Penn ID and a printout of the last web page (the receipt) to Cynthia Anderson in the OPIM Suite in Steinberg Hall-Dietrich Hall (1300 SH-DH), Monday through Friday between 9am and 4pm in order to collect your money. Funding for this study was provided by various research organizations.

All of your responses are confidential and will be reported only in aggregate and anonymously. You will not be identified by name in any publication or presentations of results from this research. If you have questions or comments about the content of the exercise, send email to crosonr@wharton.upenn.edu.

Thank you again for your participation.

Continue with Instructions

Figure 8: Screen Shot 1

The second page details the decisions subjects must make throughout the experiment, as well as the compensation scheme. Note that subjects are told that the ratio of cards may differ from deck to deck in each session.

[Webmaster email link](#)

In this exercise you will play three sessions of a game in which you will be dealt a card from one of two different decks of cards. These are not normal playing cards, but instead have either YOU WIN or YOU LOSE printed on them.

The decision you will be making will be choosing between the two decks of cards. First you will click on the picture of one of the two decks of cards to choose it. The program will then shuffle the chosen deck, randomly select one card from that deck, and show you the selected card. If the card is a WIN card, you will win 10 points. If the card is a LOSE card, you will win no points. After you have seen the card that was chosen from the deck and the points have been added to your account, the card will be returned into the deck from which it came and you will again choose between the two decks to draw your next card.

The number of times you will choose between the two decks in each session has been randomly determined. After each choice between two decks, the probability of the session ending is 1%, while with a 99% chance the session will continue, and you will be asked to choose between the same two decks again.

After the session ends by this random determination, we will change the two decks of cards to different decks. You will then play another session with the same rules and probabilities of ending and continuing. Following the second session the decks will again be replaced and a third session will be run. We will ask you to answer some questions describing your decisions between each session. After the third session ends, the exercise will be over. We will ask you to complete a post-experimental survey and to print a receipt page, which you will be able to exchange for your cash payment.

The ratio of WIN to LOSE cards may differ from deck to deck but always remains the same in the same-colored deck. While you will not be told the decks' composition, before each of the three sessions the program will show you some example cards drawn from the two decks you will use. In addition, at any time during a session you can view a history of the choices you made and their outcomes by pressing the HISTORY button located on the bottom left-hand side of your screen.

Points will accumulate over all three sessions, and the program will keep a running record of the total number of points you have earned during the exercise in your bank account shown underneath the pictures of the decks. All participants will earn a \$5 participation fee, plus \$1 for every 100 points in the account when the exercise ends. Thus the exact amount of money you earn will depend on the choices you make and the cards that appear.

Click on the button below to begin the exercise with some example draws from the first two decks.

Begin the Exercise (Example)

Figure 9: Screen Shot 2

The third page is the first screen of the actual experiment. Subjects are given three prior outcomes for each deck.

[Webmaster email link](#)

Below is a page much like the one you will face in the exercise. There are two decks for you to choose between, here the red deck and the blue deck. To choose a deck you will simply click on it or the link below it.

Underneath the decks is a box showing the result of the most recent round (WIN or LOSE). Also underneath the decks is a link to the HISTORY of your choices. Pressing this link will bring up a new window containing a history of your choices and their outcomes. You must close this window to continue with the exercise.

Underneath the history button is your bank balance. You will begin with a balance of zero. Each WIN card that is drawn earns 10 points that is added to the bank. The bank is carried over through the entire experiment.

In a moment you will be choosing between the red deck and the blue deck as depicted below.

Red Deck**Blue Deck**

Result of this Round:

[View History](#)

Current Bank Balance:

Before you start, here are some examples of cards that have been chosen at random from these decks.

<u>Red Deck</u>	<u>Blue Deck</u>
WIN	WIN
LOSE	WIN
WIN	LOSE

If you have questions, click on the first button below to review the instructions again. If not, click on the second button below to begin the first session, choosing between the red and blue decks.

Figure 10: Screen Shot 3

The fourth and fifth show what the screen looks like just before a choice is made, as well as just after the choice, when the “You Win!” or “You Lose” outcome is shown.

Figure 11: Screen Shots 4 and 5

The sixth page is a sample of the history that a subject has access to at any time during the experiment.

History of choices for Session 1:		
Resume the Game		
<u>Red Deck</u>	<u>BlueDeck</u>	
5 wins	0 wins	
5 losses	3 losses	
10 total	3 total	
<u>Round</u>	<u>Red Deck</u>	<u>BlueDeck</u>
Example	WIN	WIN
Example	LOSE	WIN
Example	WIN	LOSE
1	WIN	
2	LOSE	
3		LOSE
4		LOSE
5	LOSE	
6	WIN	
7	LOSE	
8	WIN	
9	WIN	
10	WIN	
11	LOSE	
12	LOSE	
13		LOSE
Resume the Game		

Figure 12: Screen Shot 6

Finally, the seventh page shows the message at the end of the first session along with a sample bank balance.

[Webmaster email link](#)

Choose a Deck:

Red Deck**Blue Deck**

Result of this Round:

[View History](#)

Current Bank Balance:

This is the end of the first session and was the last choice you will make between the red and blue decks. Your bank balance is 290 and will be carried over to the next session.

In a moment you will begin the second session where you will be choosing between two different decks, the green deck and the grey deck. First, we would like to ask you a few questions about your experiences in this session. Click the button below to continue.

Figure 13: Screen Shot 7

C The Expected Switching Time

In this appendix we discuss results concerning the increasing and convex properties of the expected switching time, $E[\tau]$. We sketch results for several of the paper's models.

C.1 Gittins Index Model

Chang and Lai (1987) develop approximations for the Gittins index, which are asymptotically optimal as $\alpha \rightarrow 1$. Expression (4.2) in Chang and Lai (1987) defines $E[\tau]$ for a subject that uses the approximate index. It is not difficult to show that, when properly scaled (by $\ln(1 - \alpha)^{-1}$), the expression is increasing and convex in the average quality of the current arm (e.g., Gans (2002b)).

C.2 Simple Model

Gans (2002b) applies well-known results concerning random walks to analyze an asymptotic expression for $E[\tau]$ when a subject uses the Simple model. Again, the paper's analysis shows that $E[\tau]$ is increasing and convex in the average quality of the current arm.

C.3 Myopic Model

It is not difficult to show that a common version of the Myopic index has a random walk structure that yields to the analysis in Gans (2002b). To begin, let τ be the time at which such a subject switches from arm i to arm j , and let M_0^i and M_0^j be the initial indices of the two arms.

If the computation of the index takes the form of an historical average, a form that is common when the prior distribution is a conjugate of U_s , then we have

$$\begin{aligned}\tau &= \inf \left\{ t \mid M_t^i < M_0^j \right\} \\ &= \inf \left\{ t \mid \frac{1}{r+t} \left(M_0^i + \sum_{s=1}^t U_s \right) < M_0^j \right\} \\ &= \inf \left\{ t \mid \sum_{s=1}^t (U_s - M_0^j) < rM_0^j - M_0^i \right\} \\ &= \inf \left\{ t \mid \sum_{s=1}^t X_s < b \right\}\end{aligned}$$

a random walk with increments $X_s = U_s - M_0^j$ and stopping boundary $b = rM_0^j - M_0^i$. Here r represents the number of periods' worth of data for the prior distribution M_0^i .

C.4 Hot Hand Model

The switching time for the HH- n model occurs when there are exactly n consecutive failures on the current arm. Let T be the time of the first success when sampling from the current arm, and let p be the success probability of the arm. Then

$$\begin{aligned}
\mathbb{E}[\tau] &= \sum_{t=1}^{\infty} \mathbb{P}\{T = t\} \cdot \mathbb{E}[\tau \mid T = t] \\
&= \sum_{t=1}^n \mathbb{P}\{T = t\} \cdot \mathbb{E}[\tau \mid T = t] + \mathbb{P}\{T > n\} \cdot \mathbb{E}[\tau \mid T = t] \\
&= \sum_{t=1}^n (1-p)^{t-1} p \cdot (t + \mathbb{E}[\tau]) + (1-p)^n \cdot n \\
&= \sum_{t=1}^n (1-p)^{t-1} p \cdot t + \mathbb{E}[\tau] \sum_{t=1}^n (1-p)^{t-1} p + (1-p)^n \cdot n \\
&= \frac{1}{p} (1 - (1-p)^n (np + 1)) + \mathbb{E}[\tau] (1 - (1-p)^n) + (1-p)^n \cdot n
\end{aligned}$$

Solving for $\mathbb{E}[\tau]$ and simplifying, we have

$$\mathbb{E}[\tau] = \frac{1 - (1-p)^n}{p(1-p)^n}.$$

Differentiation shows that $\mathbb{E}[\tau]$ is increasing and convex in p . We see also that $\mathbb{E}[\tau] = O\left(\left(\frac{1}{1-p}\right)^n\right)$ and is growing exponentially quickly for large n . In turn, $\frac{1}{\mathbb{E}[\tau]} = O((1-p)^n)$.

D Confidence Intervals for $\hat{\tau}$

Figure 14: Confidence Intervals for $\hat{\tau}$

E Derivation of the Index for the Simple Model

In the case of a two-armed Bernoulli bandit we have $P\{\text{good arm wins}\} = \mu_G$ and $P\{\text{bad arm wins}\} = \mu_B$. In turn, $dF_B(1) = \mu_B$, $dF_G(1) = \mu_G$, $dF_B(0) = 1 - \mu_B$ and $dF_G(0) = 1 - \mu_G$. Thus, after ω_{t-1}^i wins and ℓ_{t-1}^i losses, the index (6) equals

$$\tilde{S}(\omega_{t-1}^i, \ell_{t-1}^i) = \omega_{t-1}^i \cdot \ln\left(\frac{\mu_G}{\mu_B}\right) + \ell_{t-1}^i \cdot \ln\left(\frac{1 - \mu_G}{1 - \mu_B}\right) + \ln\left(\frac{p_0^i}{1 - p_0^i}\right).$$

Furthermore, because the prior information for the two arms is identical – two wins out of three prior trials – we assume $\ln\left(\frac{p_0^i}{1 - p_0^i}\right) = \ln\left(\frac{p_0^j}{1 - p_0^j}\right)$. Therefore, we eliminate the initial positions from both left and right arms, so that both have index

$$\tilde{S}(\omega_{t-1}^i, \ell_{t-1}^i) = \omega_{t-1}^i \cdot \ln\left(\frac{\mu_G}{\mu_B}\right) + \ell_{t-1}^i \cdot \ln\left(\frac{1 - \mu_G}{1 - \mu_B}\right)$$

at time t .

By definition $\mu_G > \mu_B$, so $\ln\left(\frac{\mu_G}{\mu_B}\right) > 0$ and $\ln\left(\frac{1 - \mu_G}{1 - \mu_B}\right) < 0$, and we define $\mathcal{U} = \ln\left(\frac{\mu_G}{\mu_B}\right)$ and $\mathcal{D} = -\ln\left(\frac{1 - \mu_G}{1 - \mu_B}\right)$. It is also not difficult to verify that, for any $\mathcal{U}, \mathcal{D} > 0$, there exist corresponding $0 < \mu_B < \mu_G < 1$. The index for arm i at time t can be written as in (10).

F BIC Scores Plotted Against Average Run Length

Figure 15: BIC Scores Plotted Against $\bar{\tau}$ for 3 Sets of Models

G BIC Scores Plotted Against Free Parameter Values

Figure 16: BIC Scores Plotted Against μ^* for the Simple Model

Figure 17: BIC Scores Plotted Against γ for the ES-1 Model

Figure 18: BIC, γ , and $ES(0)$ Values for the ES-2 Model

H Results for Probability Matching

Probability Matching models are commonly found in the behavioral literature, particularly in studies of animal behavior (e.g., Robbins and Warner (1973)).

In these studies, there typically exist two potential sources of reward, say left and right. In each trial a reward is offered in the left or the right source, but not both. Repeated trials are offered, and $P\{\text{reward in left}\}$ is i.i.d. across trials.

The Probability-Matching model hypothesizes that a subject chooses the left source with probability $P\{\text{reward in left}\}$. This behavior differs markedly from optimal behavior: to maximize expected rewards, one should choose the more likely source with probability one.

Note that the setup for this experiment differs significantly from that of ours. In a bandit problem, arms' rewards are independent of each other. Here, however, sources' rewards are perfectly negative correlated.

Furthermore, the Probability-Matching model does not carry over to bandit problems in a straightforward fashion. For example, in our experiments arms have $P\{\text{win}\}$ s of 0.15/0.4, 0.40/0.40, or 0.4/0.65, and their probabilities never add up to one, as they must for the Probability-Matching hypothesis to make sense.

Nevertheless, a number of colleagues have asked us about the performance of this type of model, so we have tested it. To fit the model to the bandit setting we normalize each arm's $P\{\text{win}\}$ by the sum of both $P\{\text{win}\}$ s. For example, in the 0.15/0.40 condition, we interpret the Probability Matching model as hypothesizing that, in a given trial, the probability that the subject chooses the arm with $P\{\text{win}\} = 0.15$ equals 0.2727 ($0.15 / (0.15 + 0.40)$).

Figure 19: Actual versus Predicted Frequency of Choosing the Left Arm

Our results show how the Probability Matching model fits our data on both an aggregate and

individual basis. Figure 19 plots subjects' observed fraction of choices against the predictions. Here, each point reflects each of a subject's 347 trials. Figure 20 displays the Probability Matching model's BIC scores, just as we have done with the other models we have tested. Note that this figure also plots analogous results for random guessing.

Figure 20: BIC Scores for Last- n and Random Choice Models

In fact, our extension of the Probability Matching model displays a poor fit to the observed data. It's performance is more akin to random guessing than to that of the other models of choice we have tested.