

Modeling Behavior in a Clinically Diagnostic Sequential Risk-Taking Task

Thomas S. Wallsten, Timothy J. Pleskac, and C. W. Lejuez
University of Maryland

This article models the cognitive processes underlying learning and sequential choice in a risk-taking task for the purposes of understanding how they occur in this moderately complex environment and how behavior in it relates to self-reported real-world risk taking. The best stochastic model assumes that participants incorrectly treat outcome probabilities as stationary, update probabilities in a Bayesian fashion, evaluate choice policies prior to rather than during responding, and maintain constant response sensitivity. The model parameter associated with subjective value of gains correlates well with external risk taking. Both the overall approach, which can be expanded as the basic paradigm is varied, and the specific results provide direction for theories of risky choice and for understanding risk taking as a public health problem.

Keywords: judgment, risk, risky choice, sequential choice, subjective probability

Cognitive psychologists and behavioral economists on the one hand and clinical psychologists on the other tend to study risk taking with different aims and in different ways. The former group, focusing on the basic processes underlying risk taking, tend to use laboratory methods and formal models as their research tools. In contrast, the latter, often with a public health focus, tend to use self-report methods to assess and, if possible, prevent risk-taking behaviors that have potential for harm or danger to the individual (Jessor, 1998; Leigh, 1999). There are exceptions in both directions, with some economists and cognitive psychologists applying their models to noneconomic contexts in more naturalistic settings (e.g., Chapman, 1996; Hamalainen, Lindstedt, & Sinkko, 2000; Wakker & Stiggelbout, 1995) and with some clinical psychologists using experimental methods to study more basic processes (Bechara, Damasio, Damasio, & Anderson, 1994; Bussemeyer & Stout, 2002; Lejuez et al., 2002), but these are the exceptions rather than the rule.

This article combines the two approaches by deriving and testing models of the decision maker (DM) in a laboratory paradigm

developed by Lejuez et al. (2002) called the Balloon Analog Risk Task (BART), which has been shown to identify individuals who are prone to highly risky behavior (Aklin, Lejuez, Zvolensky, Kahler, & Gwadz, 2005; Lejuez et al., 2002; Lejuez, Aklin, Jones, et al., 2003; Lejuez, Aklin, Zvolensky, & Pedulla, 2003; Lejuez, Simmons, Aklin, Daughters, & Dvir, 2004). This paradigm is particularly interesting for at least three reasons. First, unlike many other paradigms that are used to evaluate static cognitive decision models, this one involves sequential risk taking with feedback. To be sure, most laboratory choice and decision-making experiments involve multiple trials or repeated plays, but the models used to understand behavior in those tasks almost always assume no change on the part of the participant as he or she proceeds through the session. In contrast, modeling how the DM learns from experience is an important part of our challenge. Second, the basic BART paradigm is easily manipulated systematically, with corresponding changes in the models, in order ultimately to achieve a broad understanding of risk taking. Our present goals are more modest, but we return to this point in the Discussion section. Third, as already alluded to, analysis is at the level of the individual not the group, thereby providing a means for studying individual differences in risk taking and for relating laboratory and real-world behaviors.

Our approach mirrors in many ways the pioneering work of Bussemeyer and Stout (2002), who contrasted various models of the DM in the Bechara gambling task, which was designed to identify behavioral consequences of frontal lobe lesions (Bechara et al., 1994). Our plan in this article is first to position the present research within the broad topic of judgment, choice, and risk taking. Then, we describe the BART and briefly summarize its utility for assessing self-reported engagement in risky activities, including unprotected sex, delinquent behavior, and substance use. Next, we justify and develop multiple models of the DM's underlying cognitive processes in the BART paradigm. The following section contrasts the models using data at the level of individual DMs and looks at the relationship between estimated model parameters and reported real-world risk taking. The Discussion section, finally, considers what we have discovered about the nature

Thomas S. Wallsten, Department of Psychology and Program in Neuroscience and Cognitive Science, University of Maryland; Timothy J. Pleskac and C. W. Lejuez, Department of Psychology, University of Maryland.

Portions of this article were presented at the annual meeting of the Society of Experimental Psychologists, St. Louis, MO, March 2003; the annual meeting of the Society for Judgment and Decision Making, Vancouver, British Columbia, Canada, November 2003; the biannual conference on Subjective Probability, Utility, and Decision Making, SPUDM19, of the European Association for Decision Making, Zurich, Switzerland, August 2003; the annual Behavioral Decision Research in Management Meeting, Durham, NC, 2004; and the annual meeting of the Society for Mathematical Psychology, Ann Arbor, MI, August 2004. This research was supported by National Institutes of Health Grant R21 DA14699. We thank Jerome Bussemeyer for insightful comments on an earlier version of the article.

Correspondence concerning this article should be addressed to Thomas S. Wallsten, Department of Psychology, University of Maryland, College Park, MD 20742. E-mail: twallsten@psyc.umd.edu

of learning, option evaluation, and choice in this novel paradigm; issues that arise in testing and comparing models of the sort that we present; some open research questions, including ways to generalize the paradigm; and the benefits of translational research such as this for enhancing both basic theory and clinical practice, including improving methods of diagnosing a proclivity for engaging in risky behavior.

Judgment, Choice, and Risk Taking

Situations in which risky choice opportunities recur over time are common. Examples include engaging in lotteries or other games of chance, deciding at what time to begin the commute home, buying and selling stocks, diagnosing medical or other problems, forecasting the weather, choosing whether to use drugs or to engage in other risky behaviors, and a host of others. Generally in these cases, outcome feedback at one point in time provides information that alters the subjective event probability prior to the next opportunity. Despite the close interplay between making sequential choices and learning about event probabilities, the process is rarely studied in its entirety. Thus, most models of choice assume static or unchanging DMs, and care is taken in empirical tests to prevent feedback during relevant trials. That is so regardless of whether the models are deterministic (e.g., Kahneman & Tversky's, 1979, prospect theory and its variants) or stochastic (e.g., Regenwetter & Marley, 2001).

Notable exceptions are in the areas of probability learning (e.g., Estes, 1964; Lee, 1971) and discrimination or probabilistic categorization. Kubovy and Healy (1980) summarized and organized much of the earlier work on the latter topic. More recent research includes the combined adaptive and hill-climbing models of Busemeyer and Myung (1992) and Erev's (1998) work on reinforcement-based adjustment of decision criteria. Barron and Erev (2003) described a similar mechanism that accounts for trial-by-trial updating of a gamble's subjective value as the weighted average of the prior value and the obtained outcome.

Alternatives to static models are provided by sequential sampling (e.g., Link & Heath, 1975; Vickers, 1979) and optional stopping (e.g., Fried & Peterson, 1969; Rapoport & Tversky, 1970; Saad & Russo, 1996; Seale & Rapoport, 1997) models. However, with the exception of decision field theory and its variants (Busemeyer & Townsend, 1993; Diederich, 2003; Roe, Busemeyer, & Townsend, 2001), which build on ideas of sequential sampling, these models do not explicitly provide for learning from one trial or choice opportunity to the next. Decision field theory allows learning in terms of where the preference value starts prior to information sampling leading to a new choice.

Thus, although there are learning models for simple one-stage decisions and decision models for sequential choices, this is one of the first attempts to develop a model for the conjunction: learning to make sequential decisions. The BART paradigm introduced by Lejuez and his colleagues (2002) is ideally suited to examining the relationship between judgment, choice, and learning, particularly as they relate to risky behaviors. The paradigm was designed to mimic the characteristics of such behaviors—that is, low cost of indulgence and modest gains per incident, but with an increasing probability of severe loss. On the basis of this modeling of real-world risk contingencies, it is not surprising that performance on the BART is frequently correlated with several domains of actual

risk behavior outside of the laboratory. The present article seeks to understand the source of that success by modeling the cognitive processes underlying the behavior. The approach also provides natural directions for generalizing the research to encompass a broad range of probability, cost, and payoff conditions, thereby not only improving our understanding of risky sequential choice behavior but also our ability to predict its occurrence in individuals.

The Balloon Analog Risk Task (BART)

Lejuez et al. (2002) developed the BART to overcome difficulties inherent in self-report measures of clinical risk taking. In this computer-controlled risk-taking task, the participant is faced with a series of balloons (i.e., trials) on a computer screen (see Figure 1). For each trial, the participant sequentially clicks a button on the screen to inflate the balloon until choosing to stop or the balloon explodes. Each successful click yields a gain of $x¢$ in a temporary bank. If the participant stops pumping before the balloon explodes, the money is transferred to a permanent bank visible on the screen. If the balloon explodes following a pump, the participant loses the money in the temporary bank and earns nothing for that trial. Upon completion of the trial, the payout for the last trial and cumulative earnings to that point are shown. Then a new balloon appears and the next trial begins.

Participants are not informed of the probability structure that governs the balloons' exploding. They are told only that "it is up to you to decide how much to pump each balloon. Some of these balloons might pop after just one pump. Others might not pop until they fill the whole screen." In fact, the computer allows a maximum of n pumps, and the balloon is scheduled to explode on a random one of them with a priori probability $1/n$. Thus, the probability of its exploding on the first pump is $1/n$; on the second pump given that it did not explode on the first, $1/(n - 1)$; on the third pump given that it did not explode on the first two, $1/(n - 2)$; and so forth. In general the probability s_i that the balloon will explode on pump opportunity i given that it had not exploded on the preceding $i - 1$ trials is

$$s_i = \frac{1}{n - i + 1}. \quad (1)$$

Variables that one can manipulate in the BART paradigm include the values of x and n , the number of balloons in a sequence holding x and n constant, and the sequencing of blocks of balloons that differ in values of x and n .¹

Previous research has shown the BART to be highly reliable with adolescents and young adults both within (Aklin et al., 2005; Lejuez et al., 2002, 2004) and across (Lejuez, Aklin, Jones, et al., 2003) task administrations. Additionally, the BART has yielded significant correlations with self-reported measures of risk-related constructs including impulsivity and sensation seeking, as well as past real-world risky behaviors including drug use, delinquent activity, and risky sex in both adults (Lejuez et al., 2002, 2004; Lejuez, Aklin, Jones, et al., 2003) and adolescents (Aklin et al., 2005; Lejuez, Aklin, Zvolensky, & Pedulla, 2003). Further, in

¹ As will be indicated in the Discussion section, with slight variations in the paradigm, numerous other variables can be manipulated as well.



Figure 1. A screen shot of BART at the beginning of a balloon. Participants are told to press the “pump” button to inflate the balloon. Each successful pump increments a temporary bank by an amount of money ($\times$ cents). Participants know that the balloon will explode at some point and that when it does, they will lose the money earned on that trial. BART = Balloon Analog Risk Task.

these studies the BART accounted for significant variance in these risk-taking behaviors above and beyond that provided by demographics and self-report measures of risk-related constructs including sensation seeking and impulsivity. Typically, riskiness on the BART was indexed in terms of the average number of pumps made on balloons that did not explode.² Table 1 shows some

representative correlations between the BART index and risky behaviors.

Although the BART index is correlated with real-world risk taking, current analyses say little regarding the cognitive processes underlying risk-taking behavior in general and riskiness on the task specifically. For that purpose, it is necessary to develop a descriptive model of the processes and then to compare parameter estimates under the model to the independently assessed indices of risky behavior.

Table 1
Linear Correlations Between the BART Index and Reported Frequencies of the Indicated Behaviors

Risky behavior	Correlation
Alcohol	.28 ^a
Cigarettes	.36 ^a , .36 ^b
Polydrug use	.28 ^a , .27 ^b
Gambling	.44 ^a
Seat belt	.25 ^a
Unprotected sex	.25 ^a
Stealing	.25 ^a
Composite	.48 ^c , .42 ^d

Note. The composite is of the above behaviors in adolescents. BART = Balloon Analog Risk Task.

^a Lejuez et al., 2002. ^b Lejuez, Aklin, Jones, et al., 2003. ^c Lejuez, Aklin, Zvolensky, and Pedulla, 2003. ^d Aklin, Lejuez, Zvolensky, Kahler, and Gwadz, 2005.

Modeling Behavior on the BART

In the BART, the basic observable behavior is the DM's choice to pump the balloon or to stop at each opportunity. To explain this behavior, we developed three classes of models that assume increasing amounts of processing on the part of the DM. The simplest *baseline* model has no psychological content but provides a statistical baseline against which the other models can be compared. It assumes that DMs neither learn from trial to trial (i.e., balloon to balloon) nor choose to pump or stop in response to the current conditions. Rather, the probability that they stop at each opportunity is constant over opportunities. The remaining models, which are cognitively more plausible, have decision and response

² Balloons that exploded were excluded to avoid biasing the index.

components. The first of these, a *target* model, assumes that DMs learn from balloon to balloon but carry on no deeper evaluation. Specifically, it assumes that they set a target number of pumps prior to beginning the first balloon, which is stochastically revised upward or downward depending on whether or not the prior balloon exploded before the DM stopped pumping. The probability of stopping at each pump opportunity depends on its distance from the target. Finally, there is a family of *learning and evaluation* models that add the assumption of option evaluation to that of balloon-to-balloon learning. Representations within this family differ in the particular assumptions they make regarding how DMs update their opinions about the balloons, how DMs evaluate the options, whether they do so prior to beginning each balloon or sequentially after each pump, and how the stopping probability relates to the evaluation.

Each model from the baseline through the most complex yields a probability that the DM will pump the balloon at each opportunity. These probabilities provide the means for writing a likelihood equation for a DM's data given each model, which allows researchers to estimate and compare the various models by using maximum likelihood procedures. We describe these procedures after outlining the three classes of models, *baseline*, *target*, and *learning and evaluation*. For ease of reading, many of the formal details will be relegated to appendixes.

Baseline Model

This model assumes neither learning nor sequential decisions. Rather, the probability $r_{h,i}$ that the DM pumps on opportunity i for Balloon h is constant over all balloons and opportunities. In other words, $r_{h,i} = r$ for all h and i .³ And of course the probability that the DM stops at any opportunity is $1 - r$. The parameter r must be estimated from the data for each DM via maximum likelihood estimation (MLE) procedures.

We provide a general discussion of MLE below. However, because the estimates are easily obtained for this model, we introduce the procedure here. The likelihood of the data given a model is the product of the likelihoods of the DM's choices (assuming independence, as we do). For example, if on Balloon 1, the DM pumped 15 times and then stopped on the 16th opportunity, the likelihood of the data for that balloon under the baseline model would be $r^{15}(1 - r)$. If the DM pumped 16 times and the balloon exploded on the 16th opportunity, the likelihood of the data would be r^{16} . The likelihood of the data over all balloons under the model is the product of the likelihoods for the individual balloons. Thus, letting f denote the total number of balloons, the likelihood of the full data set is

$$L(a_1, d_1, a_2, d_2, \dots, a_f, d_f) = \prod_{h=1}^f \hat{r}^{(a_h - d_h)} (1 - \hat{r})^{d_h}, \quad (2)$$

where a_h is the number of pumps the DM took on Balloon h and d_h indicates whether the balloon exploded on the last pump ($d_h = 0$) or the DM stopped ($d_h = 1$). This is a 1-degree-of-freedom (df) model, with the single parameter $\hat{r}$. To estimate it, we find the value of r that maximizes Equation 2. It is easy to show that the expression is maximized when⁴

$$\hat{r} = \frac{\sum_{h=1}^f a_h}{\sum_{h=1}^f a_h + \sum_{h=1}^f d_h},$$

that is, $\hat{r}$ equals the total number of pumps divided by total number of opportunities (pumps plus choices to stop) over all balloons.

This model is a baseline that other, more substantive, models must outperform. The remaining models have more parameters but also much more structure representing assumed cognitive processes.

The Target Model Assuming Learning Only

One simple but adaptive approach that DMs might take to the BART is to select a target prior to beginning Balloon h , which we define as the pump opportunity, i , on which the DM is equally likely to pump as not ($r_{h,i} = .5$). The probability of pumping at the first opportunity is high, and it decreases on each subsequent opportunity such that it is .5 at the target and continues to decrease thereafter. If the DM stops before the balloon explodes, then he or she sets the target higher for the next balloon; conversely, if it explodes, then he or she sets it lower. It is reasonable to assume in this strategy that the DM alters the target by increasingly smaller amounts with experience. Ultimately, the target may remain essentially constant from balloon to balloon.

We develop the model formally in Appendix A and provide a summary of its free parameters here. t_h^* is the selected target for Balloon h . It depends on the free parameters, t_1^* , the target for Balloon 1; a_1 and a_2 , constants that control how far the target is adjusted up or down, respectively, on successive balloons; and α , which controls how rapidly the adjustment dissipates with experience. Finally, there is a free parameter, β , that controls how steeply the probability of pumping decreases as opportunity number increases (with the constraint that the probability equal .5 at opportunity t_h^*). The five parameters are estimated from the data via MLE methods, which we describe later.

A virtue of this model is that it provides a cognitively simple way for the DM to learn about and respond to the task demands. In contrast to the models in the next section, it does not posit that the DM evaluates the gains or losses of pumping, nor does it assume a Bayesian learning process.

Models Postulating Learning and Option Evaluation

The numerous models in this family all agree that for Balloon 1, the DM has a well-defined prior distribution over the probability of the balloon's exploding on pump 1. The models differ in the assumptions they make about (a) the DM's mental representation of the balloon's behavior, that is, of the stochastic process controlling the explosion probability following each pump and therefore how the DM updates his or her opinion with experience; (b) how and when the DM evaluates the outcomes of either pumping

³ We thank David Huber for suggesting this version of the baseline model.

⁴ Actually, this is the classical estimate. It is also the Bayesian estimate assuming, as is commonly done, an improper prior distribution (Hayes & Winkler, 1971, pp. 486–487).

or stopping; and (c) how response sensitivity changes with experience (i.e., how strongly the choice of whether to pump or stop is tied to outcome evaluation).

In principle, we consider 24 distinct models defined by crossing four submodels of DMs' mental representations of the balloons' behavior with three evaluation submodels with two response process submodels. Many of the specific submodels, however, can be treated quickly and dismissed as inconsistent with prior or the current data. Our presentation will proceed most smoothly by first introducing the distinct submodels, immediately eliminating those that are inconsistent with data, and finally combining those that are left into eight overall viable full models for empirical evaluation. We first consider the submodels of DMs' mental representations and how they are updated; next, the submodels of the evaluation process; then the response submodels; and finally we combine the separate parts to yield the full models of sequential decision making with learning.

DMs' Representation of the Balloons' Behavior and Updating of That Representation as Learning Proceeds

Of the four submodels that we considered for this stage, we quickly dismiss two and fully develop two others. The four submodels concur in assuming that the DM (correctly) believes that all the balloons in a series are governed by the same stochastic process. Thus, experience with past balloons informs subjective probability (SP) for the current one. They further agree that DMs update their SPs after each balloon rather than after each pump. Although this assumption may be wrong, it considerably simplifies our work, and we maintain it throughout. Finally, if this simplification is accepted, the submodels also agree that SP is updated optimally in light of the assumed structure. They disagree in their assumptions of what that structure is.

In all cases, we let $p_{h,i}$ represent the DM's SP that Balloon h will explode on pump i . It is important in all that follows not to confuse the actual probability, s_i , that a balloon will explode on pump i with the participant's estimate of that probability, $p_{h,i}$, at the time of Balloon h . The former depends on how the environment is structured and, for the present task, is given by Equation 1. The latter depends on the DM's representation of the task; his or her prior probability distribution over $p_{1,1}$, the probability that Balloon 1 will explode on Pump 1; and how he or she updates that prior with experience.

As we already indicated, the submodels differ in their assumptions about how the participant thinks that s_i changes with i . Three classes of assumptions exhaust the possibilities. They are as follows:

1. The balloon's stochastic process is nonstationary such that the probability of the balloon's exploding decreases with pump number i , that is, $s_i < s_{i-1}$ for all i . Although DMs may believe anything they wish about computer-depicted balloons, this conception seems sufficiently unlikely, and we do not pursue it further.

2. The balloon's stochastic process is stationary. Consequently, the probability of the balloon's exploding is constant over all pumps, that is, $s_i = s$ for all i . Although this, too, is not how real balloons work, it is such a simple representation that we retain it as possible and develop its implications. We call this the *stationary process submodel*.

3. The balloon's stochastic process is nonstationary such that the probability of the balloon's exploding increases with pump number i , that is, $s_i > s_{i-1}$ for all i . This assumption is consistent with BART's actual process (and with real balloons). Consequently, this assumption seems the most likely one for participants to make. Of the infinite number of ways in which DMs might imagine explosion probabilities to increase with number of pumps, we considered two.

One submodel postulates that the DM assumes that explosion probabilities approach 1 geometrically with the number of pumps. Specifically, he or she assumes that $s_i = s_{i-1} + \tau(1 - s_{i-1})$, with $0 \leq \tau \leq 1$. None of the full models that included this component survived the evaluation against data that we describe subsequently. Thus, in the interest of brevity, we do not discuss it any further. The other submodel postulates that the DM's representation of the stochastic structure is the correct one and that explosion probabilities increase with pump opportunity as specified in Equation 1. We call this the *nonstationary process, increasing probability submodel*.

We are left with two submodels of how the DM represents and updates his or her SP of Balloon h exploding on opportunity i : One submodel postulates that the DM assumes that explosion probabilities remain constant with pump; the other postulates that the DM assumes that they increase in the manner that they in fact do. We consider these in turn.

Stationary process submodel. Because under this submodel, the DM assumes that the explosion probability is constant with pump, we drop the subscript i from $p_{h,i}$ for this case and work with p_h , the DM's SP prior to Balloon h that it will explode on any pump given his or her experience with the previous $h - 1$ balloons. Under this submodel, the DM has a prior distribution over p_1 , which he or she updates in a Bayesian fashion for subsequent balloons ($h > 1$).⁵

It is easier to develop the details of the model by working with the SP that Balloon h will not explode, $q_h = 1 - p_h$, rather than with the SP that it will. Assuming that the DM treats the explosion and nonexplosion probabilities as constant over all opportunities is equivalent to assuming that he or she treats this binary event as governed by a binomial distribution with parameters currently estimated as p_h and $q_h = 1 - p_h$. It is convenient, then, to represent the prior distribution over q_1 as a beta distribution with parameters a_0 and m_0 . The more certain that the DM is about the value of q prior to observing any data, the greater is m_0 , and the greater the DM thinks q is prior to observing data, the greater is a_0 , subject to

⁵ A reviewer asked why we do not assume a cognitively simpler form of updating here and with the nonstationary model to be described next (e.g., hill-climbing procedures of some sort). The answer is that optimal Bayesian models are proving increasingly useful in understanding a wide variety of cognitive (e.g., Anderson, 1991; Shiffrin & Steyvers, 1997), perceptual (Kersten, Mamassian, & Yuille, 2004), and neural (Pouget, Dayan, & Zemel, 2003) systems. From that perspective it is parsimonious to employ them here, as well, and to use other updating methods only if the Bayesian is unsatisfactory. There is no implied assumption that the DMs consciously carry out Bayesian calculations, any more than there is when using them to understand perceptual, memorial, or neural systems; there is only the assumption that they model whatever it is that the DMs are doing. Regardless, comparing the fit of the target model to the data to that of the Bayesian learning models gives us a preliminary test of this assumption.




Figure 2. Two beta distributions differing in their estimate of, and confidence in, q_1 , the subjective probability that the balloon will not explode at Pump 1. a_0 and m_0 are the parameters of the distribution that determine its shape and location.

the constraint that $m_0 > a_0 > 0$. The mean of this prior distribution over q , that is, over q_1 is

$$E(q_1) = \frac{a_0}{m_0}. \quad (3)$$

The variance is

$$\text{var}(q_1) = \frac{a_0(m_0 - a_0)}{m_0^2(m_0 + 1)}. \quad (4)$$

Figure 2 displays an example of two beta distributions modeling two different prior opinions about q_1 . The solid-line distribution shows a higher, but less confident, estimate than does the broken-line distribution. The former has a mean of .95 and a variance of .0005; the latter, a mean of .91 and a variance of .0001.

The beta distribution is a convenient prior because it is conjugate to the binomial, which means that when the DM uses Bayes's Rule to update his or her opinion to reflect information gained by pumping the previous balloon until either choosing to stop or it explodes, the revised distribution is also beta (see, e.g., Hayes & Winkler, 1971). Further details on modeling the updating process are given in Appendix B. It is sufficient to note here that the values a_0 and m_0 are free parameters estimated from the data by means of MLE procedures that we describe subsequently.

Nonstationary process, increasing probability submodel. Recall from the earlier discussion that this submodel assumes that the DM correctly believes there is a fixed number of opportunities, n , for the balloon to explode; the a priori probability of the balloon's exploding on any pump is $1/n$; and the conditional probability of its exploding on a specific pump given that it did not explode on the previous ones is of the form shown in Equation 1. However, the DM does not know what n is. Thus, we model his or her

opinion prior to Balloon 1 with a distribution over n , which then is updated in a Bayesian fashion based on the DM's cumulating experience. This case is more complex than the previous one, because there is not an intrinsically natural way to model the prior opinion, nor are we aware of plausible distributions with conjugate conditional probability functions that we might use here. As with the prior submodel, we first discuss the prior distribution and then briefly discuss its updating, reserving the details for the appendix.

The possible values of n are the positive integers, and thus we require a discrete distribution over that domain. Our approach is to work with a discrete approximation to the gamma distribution. Specifically, we define a positive real-valued gamma-distributed random variable, n , with mean $\mu_0 = E(n)$ and variance $\sigma_0 = \text{Var}(n)$, where the subscript 0 on μ and σ reflects that the parameters govern the distribution at Point 0, prior to Balloon 1. The mean of the prior distribution is the DM's best guess of the value of n prior to any experience, and the variance indexes the DM's confidence in that guess. Both are free parameters estimated from the data by means of MLE procedures.^{6,7} Figure 3 shows two different distributions representing two different estimates of n_1 and two different levels of certainty in the estimate.

Updating the prior probability distribution over $P(n)$ via Bayes's Rule as the DM gains experience with the balloons is more complicated in this case than in the previous one; details are in

⁶ Operationally, we limit n to 1,000.

⁷ To obtain the discrete approximation to the gamma distribution, we integrate the distribution from

$$x = n - 0.5 \text{ to } x = n + 0.5 \text{ for each } n = 1, 2, \dots, \infty$$

and then normalize to account for the lost area from 0 to 0.5.




Figure 3. Two discretized gamma distributions with different means and variances. The parameters represent a decision maker's estimate of n_1 and his or her confidence in that estimate.

Appendix C. It is sufficient here to point out that the updated distribution has a new expected value, $\hat{n}$ in Equation C2 in Appendix C, which leads to the DM's updated SP of Balloon h exploding on opportunity i . The updated value, Equation C3, is reproduced here for convenience:

$$p_{hi} = \frac{1}{\hat{n}_h - i + 1}.$$

Summarizing the two learning submodels. One submodel assumes that the DM treats the balloon explosion probability as constant over pump opportunities. The DM does not know this probability but has a prior distribution over it, which is updated optimally with experience. The other submodel assumes that the DM treats the explosion probability as increasing with pump opportunities. The version that we are retaining assumes that the DM correctly represents the task structure but is uncertain as to the maximum number of pump opportunities per balloon. The DM has a prior distribution over that value, which is updated optimally with experience. Next we develop two submodels for the evaluation phase of the decision process.

Evaluation Process

Two choice models are natural candidates for inclusion here, namely, subjectively expected utility (SEU; see Luce, 2000) and prospect theory (PT; Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). Experimental evidence from behavioral decision theory suggests that PT captures unique violations of SEU theory and therefore is a better descriptive model. Consequently, we use PT to model how the DM evaluates the options of pumping and stopping. The differences between PT and SEU most relevant to us are that (a) the decision-weighting function is nonadditive, in contrast to SEU's additive SP function; (b) gains and losses are valued relative to a reference point, usually the status quo; and (c) the value function for losses is steeper than for gains. The learning submodels introduced above all yield additive SPs. Rather than

add the complexity necessary to allow them to be nonadditive, we forgo that aspect of PT and assume the identity function for the weighting function. We do, however, develop the value functions, $v(x)$, and choice rule of PT. For gains relative to the status quo ($x > 0$),

$$v(x) = x^{\gamma^+}, \quad (5)$$

and for losses relative to the status quo ($x < 0$),

$$v(x) = -\theta|x|^{\gamma^-}, \quad (6)$$

with $0 < \gamma^+, \gamma^-, \theta$. Figure 4 illustrates the value function.

Valuing outcomes as described in Equation 5 and Equation 6, the DM could plan his or her behavior in either of two ways (see also Slovic, 1966): (a) sequentially evaluate the outcomes of pumping and stopping prior to each pump opportunity or (b) determine the target number of pumps that will maximize expected gain prior to beginning each balloon. Note that under SEU theory, these two strategies coincide and yield identical predictions. However, under PT, the two yield different choice policies and predict different choice patterns. Each method in turn suggests a different response rule. We now introduce each and its respective response rule together.

Sequential evaluation and choice. Under this strategy, at each pump opportunity the DM evaluates the options of pumping and stopping. Recall that x is put in the temporary bank with each successful pump. Therefore $(i-1)x$, the amount gained prior to pump opportunity i , is the status quo for opportunity i . Thus, taking another pump yields a loss of $(i-1)x$ if the balloon explodes and a gain of x if it does not. Stopping yields no change from the status quo. Thus, the expected PT values of pumping and stopping on opportunity i are

$$E_i(\text{pump}) = (1 - p_i)x^{\gamma^+} - p_i\theta((i-1)x)^{\gamma^-} \text{ and}$$

$$E_i(\text{stop}) = 0,$$




Figure 4. The value function ($v[x]$) of prospect theory captures three characteristics of how people select among gambles: (a) Loss aversion as demonstrated by the steeper slope with losses than with gains; (b) diminishing sensitivity to increasing gains and losses; and (c) evaluations of outcomes relative to a reference point or status quo.

with p_i representing the probability of the balloon's bursting given $(i - 1)$ successful pumps. The choice rule, therefore, is to favor pumping when

$$(1 - p_i)x^{\gamma^+} - p_i\theta((i - 1)x)^{\gamma^-} > 0, \quad (7)$$

with the free parameters γ^+ , γ^- , and θ estimated from the data.

Although the output of the sequential evaluation submodel is deterministic, we believe that the choice of a response is more likely probabilistic in nature. Therefore, we require a stochastic response rule that uses the sequential evaluation (i.e., the left-hand side of Expression 7) to determine the probability of pumping or stopping. Let $b_{h,i}$ denote the left-hand side of Expression 7 and substitute $\hat{p}_{h,i}$ for p_i to incorporate an SP-updating model. We then can write the product of the evaluation as

$$b_{h,i} = (1 - \hat{p}_{h,i})x^{\gamma^+} - \hat{p}_{h,i}\theta((i - 1)x)^{\gamma^-}.$$

Note that the response model should favor pumping when $b_{h,i} > 0$, favor stopping when $b_{h,i} < 0$, and favor neither when $b_{h,i} = 0$. A logistic response model satisfies these requirements. Letting $r_{h,i}$ be the probability of the DM's pumping on the i th opportunity of the h th balloon, the model is

$$r_{h,i} = \frac{e^{\beta b_{h,i}}}{1 + e^{\beta b_{h,i}}}, \quad (8)$$

where β is a free parameter ($\beta > 0$) reflecting sensitivity of choice to the evaluation, $b_{h,i}$. Note that Equation 8 remains constant with experience. We also consider a version of the logistic model that allows β to change with experience, namely,

$$\beta = h^c, \quad (9)$$

where h indexes the balloon and c is a free parameter. Positive values of c indicate increasing sensitivity (perhaps due to greater confidence), and negative values indicate decreasing sensitivity (perhaps due to boredom or fatigue) to one's evaluation. Thus, both response submodels, Equation 8 and Equation 9, have a single free parameter.

Prior evaluation and choice. In contrast to sequentially evaluating whether or not to take the next pump, the DM may determine an optimal number of pumps prior to beginning each balloon, defining optimality in terms of the number of pumps that maximizes his or her PT-expected gains. The reference point now is the holdings (0) of the temporary bank prior to beginning the current balloon, and therefore all evaluations are in terms of gains. In general, the expected gain for i pumps on Balloon h is

$$E_h(i \text{ pumps}) = \pi_{h,i}(ix)^{\gamma^+}, \quad (10)$$

where $\pi_{h,i}$ is the probability of pumping Balloon h i times in succession without its exploding.

The number of pumps that maximizes Equation 10 depends on how $\pi_{h,i}$ is expressed, which in turn depends on the DM's model for the explosion probabilities.⁸ We denote the optimizing number of pumps as g_h . If the DM assumes that the probability of the balloon's not exploding is stationary over pumps with an estimated value of q_h for all i (Equation B1 in Appendix B), then $\pi_{h,i} = q_h^i$, and Equation 10 is optimized when

$$g_h = \frac{-\gamma^+}{\ln(q_h)}. \quad (11)$$

Alternatively, if the DM accurately assumes the nonstationary increasing nature of the balloon's explosion probability and estimates it as shown in Equation C3 in Appendix C, then $\pi_{h,i} = (\hat{n}_h - i)/\hat{n}_h$, and Equation 10 is optimized when

$$g_h = \frac{\hat{n}_h \gamma^+}{1 + \gamma^+}. \quad (12)$$

For either model, the parameter γ^+ must be estimated from the choice data.

As with the sequential case, we assume that DM probabilistically pumps or stops at opportunity i given g_h . The response model in this case uses a logistic rule constructed so that the probability of the DM's taking the i th pump on the h th balloon, $r_{h,i}$, strictly decreases with each pump opportunity and is equal to .5 when $i = g_h$. A logistic expression that provides this property is

$$r_{h,i} = \frac{1}{1 + e^{\beta \delta_{h,i}}}, \quad (13)$$

where $\delta_{h,i} = i - g_h$, and β is a response parameter representing how sensitive the DM is to his or her prior evaluation of options. The parameter β must also be estimated from the choice data. Note again that β remains constant with experience. We can incorporate experience into the response rule by means of Equation 9.

Combining the Submodels

To summarize this family of models, we considered two potential DM representations of the balloon's behavior, two methods of outcome evaluation, and two forms of translating those evaluations to choice probabilities. Taking all combinations of these submodels yields eight full models.⁹ The two submodels of the DM's view of the balloon's behavior are that the explosion probability per pump is either stationary or nonstationary, increasing with pump opportunity in the manner specified by the correct rule. In either case, the DM updates his or her probabilities according to Bayes's Rule. The two evaluation submodels both rest on prospect theory and differ in whether the evaluation is done sequentially with each pump opportunity or only prior to the first pump. Finally, the two mappings to choice probability differ in whether they are constant with experience or change with experience. The eight full learning and evaluation models, numbered 1–8, along with the baseline and target models developed previously, are summarized in Table 2.

Model Estimation and Comparison

Overview

Considering the entire development to this point, we have 10 distinct models, as summarized in Table 2. The baseline model has

⁸ Equation 10 is optimized by setting its derivative with respect to i equal to 0, solving for i , and rounding to the nearest integer. This procedure is not guaranteed to yield the optimal solution, as it treats i as continuous prior to rounding, whereas in fact i is an integer. But the solution will not be off by more than 1 and in fact the method worked well.

⁹ In fact, we considered many other submodels as well but do not present them here because they described the data so poorly.

Table 2
Summary of the 10 Models Evaluated

Model	Learning	Option evaluation	DM's model of explosion probability as pump opportunity increases	Method of learning	Response sensitivity with experience	No. of free parameters
Baseline	No	None	None	None	Constant	1
Target	Yes	None	None	Target adjustment Equation A2	Constant Equation A1	5
1	Yes	Sequential Equation 7	Stationary	Update SP Equation 4	Constant Equation 8	6
2	Yes	Sequential Equation 7	Stationary	Update SP Equation 4	Increasing Equations 8, 9	6
3	Yes	Prior Equation 11	Stationary	Update SP Equation 4	Constant Equation 13	4
4	Yes	Prior Equation 11	Stationary	Update SP Equation 4	Increasing Equations 9, 13	4
5	Yes	Sequential Equation 7	Increasing	Update SP Equation C1	Constant Equation 8	6
6	Yes	Sequential Equation 7	Increasing	Update SP Equation C1	Increasing Equations 8, 9	6
7	Yes	Prior Equation 12	Increasing	Update SP Equation C1	Constant Equation 13	4
8	Yes	Prior Equation 12	Increasing	Update SP Equation C1	Increasing Equations 9, 13	4

Note. DM = decision maker; SP = subjective probability.

no structure and one free parameter. The remaining models have substantial cognitive structure and four to six free parameters each. The target model assumes simple balloon-to-balloon learning in the form of target adjustment but no outcome evaluation. The remaining models (Models 1–8) all assume both balloon-to-balloon learning and outcome evaluation, but they differ in various important ways, as described earlier.

Some Details on Model Estimation

Each of the models was estimated separately for each DM by means of the MLE method. This section provides a brief overview of the approach, with details relegated to Appendix D. Recall from the previous presentation of the baseline model that the likelihood of the data over all pumps and balloons under a given model (assuming independence) is the product of the model probabilities of each response (pumping or stopping). The expression (Equation 2) is easy to write in the case of that model because its probabilities are assumed constant over all balloons and opportunities per balloon. That is not the case with the remaining models (i.e., the target and Models 1–8).

The more complicated version of Equation 2 is provided in Appendix D as Equation D1. The important point for the moment is that Equation D1 shows the likelihood as a function of the probabilities, $r_{h,i}$, of pumping Balloon h on opportunity i over all balloons and opportunities. Each model obtains these probabilities in a different way. Thus, Equation D1 is completed for a given

model by substituting the model's expressions for the $r_{h,i}$, which can be constructed by combining across the appropriate submodels following the guide in Table 2. The expressions involve free parameters (e.g., four in the case of Model 3), for which estimates are sought that maximize Equation D1. Additional details are in Appendix D.

The Data Used to Compare the Models

To fit and compare the models, we used data from Lejuez, Aklin, Jones, et al. (2003). The purpose of that study was to examine differences in behavior on the BART for smokers and nonsmokers in order to test the hypothesis that smokers show higher levels of riskiness than nonsmokers. Individuals were recruited into the study without restriction if they were regular smokers ($\geq$ five cigarettes/day for at least 6 months) or nonsmokers, resulting in a sample of 34 nonsmokers and 24 daily smokers. Smoking status was balanced across gender, and participant age ranged from 18 to 30 years ($M = 20.1$, $SD = 2.8$). Of the total, 68% were Caucasian, 13% were Asian American, and 12% were African American (the remaining 7% marked "other" on the demographics form).

Sessions lasted approximately 1 hr, during which participants completed three 30-balloon blocks of the BART, with other tasks occurring between the blocks. The maximum number of pumps was set at 128 for each balloon. Therefore the a priori probability of a balloon's exploding on any given pump was 1/128, and the

conditional probability of each one exploding on Pump i given that it had not exploded on the previous pumps was (from Equation 1) $s_i = 1/(129 - i)$.¹⁰ The DMs accumulated \$.05 in their temporary bank for each successful pump. The winnings from each balloon were moved to a permanent bank if the DM stopped pumping before the balloon exploded and were lost otherwise.

The tasks intervening between the blocks of BART trials included another decision-making task and the completion of self-report questionnaires assessing smoking status and other risk behaviors, including use of other drugs, as well as personality characteristics, including levels of sensation seeking and impulsivity. At the start of the session, participants were informed that they could make between \$10 and \$30, with the exact amount based on the comparison of their total scores for both decision tasks and the total scores of the other participants.

Results from the study indicated that smokers were more risky on the BART than were nonsmokers across each of the three administrations of the task. Further, BART scores were correlated with the number of other drug classes tried in the past year across the following categories: (a) alcohol, (b) cannabis, (c) nicotine, (d) cocaine, (e) stimulants, (f) sedatives, (g) opiates, (h) hallucinogens, (i) designer drugs, and (j) other. Some of these results are included in Table 1.

We used only the first session of 30 balloons in order to fit the models.¹¹ The adjusted BART score (the average number of pumps taken on balloons that did not explode) for these 30 trials was 37.6, with a standard error of 1.5. Balloons exploded on 34.1% of the trials. Figure 5 provides a convenient means of visualizing the general structure of the data. It shows the mean probability of pumping, $\hat{r}_i$, at each of the 128 opportunities across all 58 participants, with 95% confidence intervals. As the figure shows, the probability of pumping decreased with opportunity in an ogival fashion, but there was a large amount of individual variability. We fit all models separately to the data of each participant, not to the average results.




Figure 5. The average probability of pumping for each of the 128 opportunities across all 58 participants, obtained from the estimates used in the baseline model. The dotted lines represent the 95% confidence interval for each pump's estimate. Prob. = probability.

Model Comparison and Selection

Before continuing further we should emphasize that because the models are not nested, standard maximum likelihood ratio tests are not available to evaluate them. One may still compare maximum log likelihoods (MLL; see Appendix D) of the data under competing models at a descriptive level, but in doing so one must worry about differential model complexity, which affects the flexibility of a model to fit disparate data patterns (Myung & Pitt, 2002; Pitt, Kim, & Myung, 2003; Pitt, Myung, & Zhang, 2002; Roberts & Pashler, 2000). Pitt, Myung, and Zhang (2002) contended that model comparison should be based on analyzing models' response surfaces in order to adequately comprehend their complexity, with an index such as minimum description length used as a means of model selection. This method is difficult if not impossible with the present models because of their very complex forms and is beyond the scope of this article. Future work will pursue the consequences of these ideas on the present model or on related ones.

We used two methods to evaluate and select the most descriptive model. The first was to use the often-invoked Akaike information criterion (AIC; Akaike, 1973), which penalizes models for one aspect of their complexity: number of parameters. The criterion equation that combines both a model's MLL and its number of parameters, k , is

$$AIC = -2MLL + 2k.$$

Lower values are better. The second method was to compare the surviving model(s) in terms of correlations between their free parameters and external risk indices of the sort shown in Table 1. Valid models in this task ought to yield parameters that correlate with these indices.

Table 3, in which the data of each participant are treated separately, shows how many participants were best fit by each model using AIC. In addition, the table shows the mean AIC and mean MLL of each model. Note that almost all the formulations assuming learning (SP updating) and PT evaluation of the options do substantially better than either the baseline or the target models (Model 4 is an exception). Recall that the baseline model assumes that DMs neither learn nor evaluate options. Clearly, that is wrong. The target model does assume learning in the form of balloon-to-balloon adjustment of a target number of pumps but not complex

¹⁰ The optimal expected-value strategy assuming full knowledge of the task structure was to pump 64 times and then stop. Explosion points were determined for each balloon in the manner described (i.e., each pump had an a priori probability of 1/128 of yielding an explosion) but with the constraint that explosions were scheduled to occur on average on Pump 64 over the entire 90 balloons and within each subblock of 10. In addition, the first balloon in each block of 30 was scheduled to explode on Pump 64 to avoid biasing the DM to be either risk prone or risk averse. Although these constraints were useful from the point of view of developing the BART as a diagnostic test, it would have been better from the perspective of model evaluation to determine the explosion point without them. We have no reason to believe, however, that this procedure affected the overall pattern of results.

¹¹ We also tested the models on all 90 trials for some participants. The results were unchanged, but the computer results took so long that we concentrated on the first 30 trials for the extensive model comparisons.

Table 3
Model Comparisons Showing the Number of Free Parameters per Model (df), the Number of DMs Best Fit by Each Model According to AIC, and the Mean Values of the Two Fit Statistics (AIC and MLL)

Model	df	No. of DMs best fit	Mean AIC	Mean MLL
Baseline	1	0	192.79	-95.40
Target	5	5	164.39	-77.20
1	6	15	154.96	-71.48
2	6	9	156.46	-72.23
3	4	16	154.14	-73.07
4	4	4	186.97	-89.49
5	6	1	160.28	-74.14
6	6	1	161.35	-74.67
7	4	5	159.08	-75.54
8	4	2	158.96	-75.48

Note. DM = decision maker; AIC = Akaike information criterion; MLL = maximum log likelihood.

cognitive processing beyond target adjustment. Clearly, that, too, is wrong. By firmly rejecting these models in favor of the others, we can conclude that the DMs learn something about the probabilistic structure of the situation and balance outcomes against probabilities in making their choices.

For more fine-grained conclusions, note that two of the eight evaluation models, Models 1 and 3, fit a majority of the participants best. These two models assume that (a) DMs inaccurately treated the balloon's stochastic process as stationary but optimally updated their opinions about the parameters of the process and (b) response sensitivity (or response variability) remained constant across balloons. Models 1 and 3 differ in that the former assumes that DMs sequentially evaluate their options at each pump opportunity, whereas the latter assumes a single evaluation prior to beginning to pump the balloon.¹² On the basis of both AIC and the number of DMs best fit, Model 3 is marginally better than Model 1; we therefore summarize its MLE parameters in Table 4 across DMs.

External Validity

Riefer, Knapp, Batchelder, Bamber, and Manifold (2002) pointed out that parameter estimates of models that fit the data well do not necessarily reflect the underlying psychological constructs, and therefore it is useful to perform some sort of external validation. One method for doing so is to experimentally manipulate conditions that should change parameter values in predictable ways. Another method, which we use, is to correlate the parameter estimates with measures of independently recorded behaviors they are meant to reflect. As the adjusted BART scores are intended to (and do) correlate with risky behaviors of the sort listed in Table 1, the model we select as best should also provide output that correlates with these behaviors.

Recall that participants completed a battery of self-report scales, including questions about how often in the past year they had unprotected sex, rode in the passenger seat of a car without a seatbelt, engaged in reckless driving, and stole an item. Questions were asked as well about their use of nicotine and 9 other drug categories (cannabis, cocaine, MDMA [ecstasy], stimulants [e.g.,

speed], sedatives/hypnotics, opiates, hallucinogens, phencyclidine [PCP], and inhalants), for a total of 10 drug categories. For present purposes, we summed the number of drugs endorsed to yield an index from 0 to 10. Because the self-reports at most rank-order the participants according to risk propensity in any given category, we used Goodman and Kruskal's G (see Gonzalez & Nelson, 1996) to assess the ordinal association between riskiness in that category and the BART data or model parameter estimates, limiting attention to the two most successful models (Models 1 and 3).

Table 5 shows the results. The first row provides the ordinal correlation (Goodman & Kruskal's G) between the adjusted BART score for the 30 trials used in our model analyses and the various risk categories over the 58 participants. This association is significant and positive for all the categories except seatbelt use, essentially paralleling results with other BART studies (cf. Table 1). The remaining rows of the table are partitioned into three panels, one for each of the surviving models, 3 and 1, and a final panel to which we return in the Discussion section. Row 1 of the first two panels concerns an interesting summary statistic, g_{31} , representing the output of each model. To understand g_{31} , recall that for a given set of parameter estimates, each model yields a probability distribution over pumping and stopping per opportunity per balloon; g_{31} is the pump opportunity on Balloon 31, after the DM has experienced all 30 actual balloons, for which under the model with its optimal parameter estimates the DM has a .50 probability of stopping. If a model is valid, g_{31} should correlate with the external risk indices roughly as well as does the adjusted BART score. Note that the pattern of correlations for g_{31} under Model 3 is similar to that for the adjusted BART, whereas that for g_{31} under Model 1 is weaker.

The correlations of the MLE parameter estimates for each model are in the remaining rows of the Model 1 and Model 3 panels. The bottom four rows of each panel show correlations for two forms of the parameters summarizing the DMs' prior opinions, a_0 and m_0 on the one hand and $E(q_1)$ and $\ln(\text{var}(\hat{q}_1))$ on the other (see Equations 3 and 4). Note that for Model 3, the pattern of correlations for the value function exponent, γ^+ , mirrors that for g_{31} . In addition, the response-sensitivity parameter, β , is negatively correlated with the number of drug categories. In contrast, none of the Model 1 parameters correlates with the external risk categories. This overall pattern of results suggests, but certainly does not prove, that Model 3 is a better descriptor than Model 1. Cautions about and possible interpretations of this pattern of correlations are best deferred to the Discussion section.

Focusing attention only on Model 3 now, we illustrate that the model accounts for a large amount of the variance in the data, both within and between participants. Consider first within-DM balloon-to-balloon variance. Figure 6 plots the model's predicted probability of pumping for Participant 11 on his or her first four balloons (1–4) in the top panel, middle four balloons (13–16) in the center panel, and last four balloons (27–30) in the bottom

¹² We also tried versions of these two models that included a parameter for response bias in addition to the one for response sensitivity used in Equation 8 and Equation 13, but the results were not substantially or significantly improved.

Table 4
Maximum Likelihood Parameter Estimates for Model 3

Statistic	β	γ^+	a_0	m_0	$E(q_1)$	$\text{Var}(\hat{q}_1)$
Mean	.15	0.81	14,024.59	14,064.67	.99	6.76E-04
First quartile	.07	0.45	35.46	36.04	.98	2.78E-05
Median	.12	0.70	119.92	121.43	.99	8.71E-05
Third quartile	.19	1.05	296.75	299.62	1.00	3.11E-04
IQR	.12	0.60	261.29	263.57	.01	2.83E-04

Note. IQR = interquartile range; E denotes expected value.

panel. The probabilities are calculated by using the MLE estimated parameter values for this participant and the specific conditions imposed by the behavior of the DM and the prior balloons. Graphs for balloons that ultimately exploded are plotted with open circles; those for balloons on which the DM stopped are plotted as solid lines.

Notice first the large difference in the predictions for the first four balloons, as compared with the later balloons, for this DM. He or she took many pumps on the early balloons and then, with learning and experience, became more likely to stop after fewer pumps. The model also captures sequential effects in behavior. Observe Model 3's prediction for Participant 11's third and fourth balloons. The model shows that after the participant's first two

balloons, on which he or she stopped successfully without the balloon's exploding, the DM was likely to increase the number of pumps on Balloon 3. Unfortunately, that one exploded after 39 pumps, causing the model to predict many fewer pumps on Balloon 4. In fact, the DM stopped after 20 pumps. The same general pattern is exemplified in the model's predictions for subsequent balloons. But the later balloons also show asymptotic learning in that explosions and successful stopping begin to have less and less of an effect on the DM's performance. Corresponding plots for the other DMs similarly reflect their parameter estimates and unique trial-by-trial experiences.

Now consider between-DM variance. Figure 7 shows Model 3's predictions for four participants (51, 46, 3, and 7) selected to illustrate participants who differed substantially in their BART as well as in their self-reported risky behaviors. The first four rows of Table 6 show the number of drug categories reported by each participant as well as their optimal Model 3 parameters. The panels of Figure 7, ordered according to number of reported drug categories by each participant, show the model's predicted probabilities for the first two (1 and 2) and the last two (29 and 30) balloons per DM. The large variability in the model's predictions for each participant is apparent when one scans the four panels. Note that the nature of the change in balloon-to-balloon probabilities varies substantially over participants and that by Balloon 30, the probabilities reflect the ordering by drug categories well. (Although the

Table 5
Ordinal Correlations (*G*) Between the Indicated Risk Categories and the Adjusted BART Score, and Between the Categories and Various Model Characteristics

Index	Drug categories tried ^a	Unprotected sex	Seatbelt use	Reckless driving	Stealing
Adjusted BART score	.332*	.370*	.079	.281*	.273*
Model 3					
g_{31}	.321*	.313*	-.003	.127	.391*
β	-.204*	-.151	.103	-.034	-.302
γ^+	.287*	.307*	.025	.095	.456*
a_0	-.138	-.246	.003	-.036	-.232
m_0	-.136	-.248	.008	-.036	-.227
$E(q_1)$	-.139	-.150	-.170	.112	-.120
$\ln(\text{var}(\hat{q}_1))$	.160	.259	.009	-.016	.191
Model 1					
g_{31}	.191*	.173	-.061	-.024	.465*
β	.117	.158	.041	.042	-.077
$\gamma^+ - \gamma^-$	-.028	-.237	-.171	-.132	.222
$1/\theta$	-.022	.206	.098	-.005	-.068
a_0	.016	-.248	.068	.077	-.060
m_0	.016	-.248	.068	.077	-.063
$E(q_1)$	-.063	-.041	-.192	.045	-.122
$\ln(\text{var}(\hat{q}_1))$	.095	.127	.043	.001	.096
Additional results					
Best-fit model g_{31}	.280*	.263*	.148*	.127*	.270
Baseline $\hat{r}$	.341*	.373*	.046	.279*	.373*

Note. The adjusted Balloon Analog Risk Task (BART) score is an observed variable; β , γ^+ , $1/\theta$, a_0 , and m_0 are all estimated model parameters; and g_{31} , $\gamma^+ - \gamma^-$, $E(q_1)$, and $\ln(\text{var}(q_1))$ are all derived from estimated model parameters.

^a 10 possible.

* $p < .05$.




Figure 6. Model 3 predictions for Participant 11's first four balloons, middle four balloons, and last four balloons. The numbers $x:y$ associated with each curve represent balloon number (x) and number of pumps taken (y). Prob. = probability.

ordering is perfect among these four DMs, it is not over the whole sample, as illustrated by the .321 correlation shown in Table 5 between g_{31} and number of drug categories tried.) Model 3 also captures the strategy common to many participants of sacrificing

the first balloon to learn about the environment. Three of the four DMs illustrated here pumped the balloon until it exploded on Pump 64 (see Footnote 10), whereas one DM did not, and that is well reflected by the model for these DMs as well as for the others.




Figure 7. Model 3 predictions for 4 participants' first and last two balloons. The numbers $x:y$ associated with each curve represent balloon number (x) and number of pumps taken (y). Prob. = probability.

Table 6
*Number of Drugs Tried and Model 3 Parameters for the
 Participants Illustrated in Figures 6 and 7*

Participant	No. of drug classes tried	β	γ^+	a_0	m_0
51	9/10	.034	1.81	35.37	35.82
46	8/10	.25	0.44	76.83	77.14
3	1/10	.22	0.03	28,188.73	28,202.90
7	0/10	.18	0.16	256.15	256.17
11	1/10	.13	0.56	80.93	81.93

Discussion

It is useful to break this section into five parts. The first part provides an overview of developments and results to this point; the second interprets and discusses Model 3, the model most consistent with the data; the third focuses on some of the problems we encountered in model estimation and evaluation; the fourth discusses open issues and future directions, including ways in which to generalize the paradigm; and the final part discusses the benefits of translational research such as this for enhancing both basic research and clinical practice.

Overview

We began with a sequential decision-making task, the BART, designed and validated for the purpose of identifying individuals prone to excessive risk taking. Our aim was to interpret behavior in this task in terms of its underlying cognitive processes, for the dual purposes of understanding sequential choice in this moderately complex environment and understanding why the adjusted BART score correlates with self-reported real-world risk-taking.

To these ends, we proposed increasingly complex models of the DMs' choice behavior. Simplest was the baseline model, which assumed that DMs neither learned from balloon to balloon nor engaged in any sort of outcome evaluation. Rather, DMs simply pumped with a fixed probability. Next in complexity was a target model, which assumed a simple form of learning from experience but no outcome evaluation. In this conception, DMs set a target prior to each balloon that governed the pumping probabilities. Target values were adjusted down or up from balloon to balloon according to whether or not the previous balloon had exploded, with adjustments decreasing in magnitude as the DM gained experience in the task. The remaining models all assumed more complex cognitive processing involving both optimal learning and option evaluation. We rejected various models in this family early on but maintained eight (Models 1–8) for thorough evaluation.

These eight models all agreed that options were evaluated as described in prospect theory, but they differed orthogonally on three classes of assumptions. The DMs evaluated whether to pump or stop either prior to beginning a balloon or sequentially with each pump. The models all agreed that DMs optimally updated their opinions about the explosion probabilities from balloon to balloon but differed in terms of whether the DM's mental model included stationary or nonstationary increasing explosion probabilities over pumps (with the form of the increase governed by the rule that actually was in effect). Finally, the models differed in terms of

whether response sensitivity was constant or increased over balloons.

The data clearly contradicted the baseline and target models in favor of the others, thereby demonstrating that the DMs were both learning about the environment and evaluating options. Of the remaining models, 1 and 3 had the best average goodness-of-fit measures and were the best-fitting models for the greatest number of participants. Ultimately, we favored Model 3 over Model 1 because it had slightly better goodness-of-fit statistics, and more important, because it correlated better with the self-reported risk behaviors.

Model 3

Recall that this model assumes that prior to beginning each balloon, DMs engage in a prospect-theory-like evaluation of the expected outcomes to determine an optimal number of pumps. Their pumping probability then decreases with opportunity number such that it equals .5 at the optimal stopping point. Model fit was not improved either by including a bias parameter that allowed the pumping probability at the goal to differ from .5 or by allowing the response sensitivity parameter to change with experience. The remaining feature of Model 3 was its assumption that the DM incorrectly treated the conditional explosion probability as constant rather than as increasing over pumps.

In terms of cognitive economy, the strategy represented by this model is not nearly so easy to implement as those embodied in the baseline or target models but requires much less mental computation than the alternative strategies represented by the others. From the perspective of fast and frugal heuristics (Gigerenzer & Todd, 1999), Model 3 may be a good compromise between the simple but very maladaptive baseline and target strategies and the more complex strategies implied by the other models.

It is notable that the summary statistic g_{31} for Model 3 correlates with the self-reported risky behaviors to the same extent that the adjusted BART score does. The additional finding that γ^+ (the prospect-theory value-function exponent) is the only parameter to reflect the g_{31} correlations provides an explanation for BART's success. It suggests that the primary factor affecting the DMs' risk propensity in both the BART task and the external risk categories is how they value outcomes. Specifically, risky individuals in these tasks tend to be those who marginally discount gains the least. In fact, Model 3 provided the best fit to the data of some DMs when γ^+ was allowed to go above 1, suggesting that those DMs inflate the marginal value of gains. These same individuals also tended to report the largest number of risky behaviors. An alternative explanation for risky behavior would be that DMs differ in their initial opinions of the probabilities of the negative outcomes and that risk-seeking DMs have lower prior subjective explosion probabilities than do DMs who are risk averse. In contradiction to that view, however, the parameters reflecting subjective probabilities did not correlate with the self-reports.

This pattern of results is consistent with that of Stout, Busemeyer, Lin, Grant, and Bonson (2004), who used the Busemeyer and Stout (2002) expectancy-valence model to contrast the behavior of cocaine abusers and a "normal" control group on the Bechara gambling task (Bechara et al., 1994). That model has three parameters: a motivational one reflecting amount of attention given to gains versus losses, a parameter reflecting trial-to-trial learning, and a third reflecting response consistency. They found

that only the motivation parameter, which corresponds roughly to our γ^+ , distinguished the cocaine abusers from the controls. There was no difference in the learning parameters, which correspond roughly to our SP parameters, nor in the response consistency parameters.

In the present study, the response sensitivity parameter β was also significantly negatively associated with a single risk category: the composite drug score. It may be that individuals who use illicit drugs are less responsive to their outcome evaluations than are individuals who do not, but this may also just be a chance result that must be replicated before taken too seriously.

Problems in Model Estimation and Evaluation

As anyone working with complex models of the sort treated here knows, finding optimal solutions is not a straightforward matter. To begin, local maxima in the log-likelihood function, Equation D1 in Appendix D, abound, making finding the global maximum difficult. To mitigate this problem, we used a combination of a grid search technique and the Nelder–Mead method, as described in Appendix D, to estimate each model described in the model estimation section. However, this is not a perfect method, and better ones surely exist. A potentially better approach would be to specify a joint probability distribution over the parameter space and to sample starting values accordingly, but the computational issues are problematic. Other difficulties in formal model comparison arise from the non-nested nature of the models, as discussed earlier.

We must emphasize the different degrees of confidence that one should have in rejecting some models, and the cognitive processing they imply, in favor of other models and implied processes. We are very confident in rejecting the baseline and target models, which imply minimal cognitive processing, because all of the alternatives do better. Selecting among the remaining models is somewhat more tenuous, and opinion should remain open about them. Although Model 3 seems best given the present paradigm and data, it remains to be determined whether it will survive additional scrutiny. For example, the conclusion that the DMs made prior rather than sequential evaluations was a close call, and subsequent research may reverse it. The conclusion of no response bias and of constant response sensitivity may have depended on how we expressed those factors. We selected a response function that is commonly used because of its mathematical convenience, not because of its psychological realism. Other response functions may lead to different conclusions. Finally, the conclusion that DMs misunderstood the explosion process as stationary rather than as non-stationary-increasing may change with different perceptual or structural representations of the process.¹³ Indeed, Pleskac (2004) showed in a related paradigm that DMs correctly represent dynamic event probabilities as stationary or increasing when the processes are openly visible and do not have to be learned with experience.

Open Issues, Future Directions, and Paradigm Generalization

One important and vexing question is how to represent individual differences. It is apparent from Table 3 that a single model was not uniformly best for all DMs. Nevertheless, we used Model 3 for everyone when correlating estimated parameters with external risk

scales and interpreting the results. This approach is reasonable on the assumption that cognitive processing is fundamentally similar in all DMs and that a single model would in fact fit everyone best were it not for statistical and computational issues. Instead, because of limited data and problems that arise in optimizing functions, other models appear to be better for some individuals. In this view, individual differences, due, for example, to different life experiences, levels of motivation, and so forth, are reflected in the parameter estimates of a single best model, which can be selected only when looking at the full pattern of results across many individuals. Batchelder and colleagues (Batchelder, 1998; Riefer et al., 2002) use the term “cognitive psychometrics” to describe this approach. Although this is the approach that we have taken, an equally reasonable alternative perspective is that individual differences are manifested in the very nature of the cognitive processing, and therefore different models and different parameters within models are appropriate across DMs.

Fully contrasting these two approaches is beyond the scope of the present article. However, as a first step, we used each DM’s best model to estimate a common derived parameter, which we then correlated with the external risk indices. That parameter, termed g_{31} when estimated for Model 3, is the pump opportunity for Balloon 31 at which the stopping probability is .5. Recall from Table 5 that g_{31} from Model 3 correlated well with the external indices. If the multiple-model approach is superior to the single-model approach, then the correlations obtained using values of g_{31} estimated from the individually best models should equal or exceed those obtained with values from Model 3 only. The results are in the penultimate row of Table 5, labeled *Best-fit model g_{31}* , and provide little comfort to the multiple-model proponents.

This outcome does not settle the issue of whether individual differences in risky sequential choice are best understood in terms of distinct cognitive representations (multiple models) or as variations within a single representation (one model with distributions of parameter values). But it does provide a methodology for investigating it and a first cut at an answer.

The last row of Table 5 shows the results when correlating the pump probability, $\hat{r}$, estimated from the baseline model (which, of course, is that model’s g_{31}) with the external risk indices. They are comparable to (indeed, very slightly better than) those obtained with g_{31} from Model 3 or with the adjusted BART score, and they are much better than those obtained with the individually best-fitting models. The reason the baseline model does so well is that it is merely a summary of the data without imposing any constraints on them. Only Model 3 imposes the proper constraints, which is further evidence against the multiple-model approach.

The next question is, if the baseline model predicts external risky behavior so well, why is any other needed? The answer is that although the baseline model predicts external behavior at least as well as the adjusted BART score does, it provides no insight into the determinants of the behavior. Model 3, in contrast, provides a reasonably deep theoretical understanding and ultimately may be useful for clinical purposes.

¹³ It is unlikely that the DMs’ error is due to insufficient sampling (30 trials), as the model conclusions did not change when we included responses to all 90 balloons for some DMs. (See Footnote 10.)

We hasten to point out that our conclusions must be evaluated and, where necessary, qualified and modified as other paradigms are studied. This research began adventitiously, taking advantage of the BART's laboratory properties and the fact that it had already been validated as a good predictor of real-world risky behavior among the populations in which it has been evaluated. However, studying behavior on the BART alone is not sufficient for gaining a full understanding of sequential risky choice. For example, it is necessary to contrast behavior when events probabilities are already known or perceptually evident versus when they must be learned, as is the case in the BART. In fact, Pleskac (2004) developed an alternative sequential risky choice task with this principle in mind. Results with it show that the BART's procedure of concealing the probabilistic structure reduces its correlation with risky behaviors. This new task also allowed simultaneous estimation of both the weighting and the value function in prospect theory.

Pleskac (2004) has made a good start, but much more remains to be done. It is necessary to explore behavior in response to a wide variety of event probability structures and under a wide variety of costs and payoffs. Researchers will achieve a full understanding only by systematically varying the paradigm and determining corresponding changes in underlying cognitive processes.

Translational Research

The present study is an example of translational research, the melding together of clinical and basic cognitive approaches in a manner that enhances both streams of work (National Institute of Mental Health, 2000; Onken & Bootzen, 1998; Zvolensky, Lejuez, Stuart, & Curtin, 2001). Translational research is especially timely in the domain of risk taking given the substantial empirical and theoretical advances within, and lack of integration between, distinct fields. In this spirit, the conclusions from the present research are provocative and require follow-up. As new paradigms are studied and we learn which processes are common and which are unique to specific types of sequential risky choice, prediction and treatment may be improved because instead of simply describing the occurrence of risky behavior, we will understand the factors underlying them. Such precision also will bear upon our current conceptualizations of risky behavior and the extent to which these behaviors are orthogonal or present along related dimensions (cf. Weber, Blais, & Betz, 2002).

In addition to increasing our basic understanding of risk taking and risk perception, research along these lines may shed light on specific clinical disorders. In this regard, a large literature on expectancies has emerged examining positive and negative reasons for substance use. Although responsible for several important insights into motives for engaging in risky substance use (Leigh, 1989), this research could greatly benefit from the level of precision that is provided with the modeling approach presented here in terms of identifying to what extent substance use is related to deficits in understanding risk versus oversensitivity to potential gains from risky behavior. As another example, Ladouceur, Sylva, Letarte, Giroux, and Jacques (1998) suggested that pathological gamblers misunderstand probability. Although that may be true (most people misunderstand probability), our conclusion that risky behavior is mitigated primarily by the DMs' discount/inflation factors for gains rather than by their perceptions of the

probabilities suggests that it is an insufficient explanation. Is the discrepancy between the two conclusions real, perhaps associated with different populations, or is it due to different methods of analysis? (Erev, Budescu, & Wallsten, 1994, provided another example of a major discrepancy in the literature that was traced primarily to differences in methods of data analysis.)

It is noteworthy that using a single model, Busemeyer and Stout (2002) found a different pattern of parameter values when contrasting the behavior of patients with Huntington's disease to that of controls than Stout et al. (2004) found when contrasting the behavior of cocaine abusers to that of controls. The Huntington group differed from the controls in the learning parameter only, whereas the cocaine abusers differed in the motivational parameter only, a result that corresponds to behavioral and neurologic differences in the groups. It is important, therefore, not to interpret our current results too broadly but rather to take them as a beginning step.

In terms of clinical and public health relevance, the current research is promising because it moves beyond merely describing risk taking and provides direction for understanding the factors that underlie such behavior. Indeed, the ability to determine the cause of risk taking on a case-by-case basis can have great implications for developing and implementing individualized risk-taking prevention programs (Botvin, 2004; Kelly & Kalichman, 2002). Although far from achieving this goal at the present time, this approach begins to combine cognitive and clinical considerations in a way that may ultimately allow us to understand and address risk-taking, both as a concept and as a public health problem.

References

- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrox & F. Caski (Eds.), *Second International Symposium on Information Theory* (pp. 267–281). Budapest, Hungary: Akademiai Kiado.
- Aklin, W. M., Lejuez, C. W., Zvolensky, M. J., Kahler, C. W., & Gwadz, M. (2005). Evaluation of behavioral measures of risk taking propensity with inner city adolescents. *Behaviour Research and Therapy*, 43, 215–228.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98, 409–429.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16, 215–233.
- Batchelder, W. H. (1998). Multinomial processing tree models and psychological assessment. *Psychological Assessment*, 10, 331–344.
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50, 7–15.
- Botvin, G. J. (2004). Advancing prevention science and practice: Challenges, critical issues, and future directions. *Prevention Science*, 5, 69–72.
- Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General*, 121, 177–194.
- Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, 14, 253–262.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100, 432–459.
- Chapman, G. (1996). Temporal discounting and utility for health and money. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 1–17.

- Diederich, A. (2003). MDFT account of decision making under time pressure. *Psychonomic Bulletin & Review*, 10, 157–166.
- Erev, I. (1998). Signal detection by human observers: A cutoff reinforcement learning model of categorization decisions under uncertainty. *Psychological Review*, 105, 280–298.
- Erev, I., Budescu, D. V., & Wallsten, T. S. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, 101, 519–527.
- Estes, W. K. (1964). Probability learning. In A. W. Melton (Ed.), *Categories of human learning* (pp. 89–128). New York: Academic Press.
- Fried, L. S., & Peterson, C. R. (1969). Information seeking: Optional versus fixed stopping. *Journal of Experimental Psychology*, 80, 525–529.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Gonzalez, R., & Nelson, T. O. (1996). Measuring ordinal associations in situations that contain tied scores. *Psychological Bulletin*, 119, 159–165.
- Hamalainen, R. P., Lindstedt, M. R., & Sinkko, K. (2000). Multiattribute risk analysis in nuclear emergency management. *Risk Analysis*, 20, 455–467.
- Hayes, W. L., & Winkler, R. L. (1971). *Statistics: Probability, inference, and decision*. New York: Holt, Rinehart & Winston.
- Jessor, R. (1998). New perspectives on adolescent risk behavior. In R. Jessor (Ed.), *New perspectives on adolescent risk behavior* (pp. 1–10). Cambridge, England: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- Kelly, J. A., & Kalichman, S. C. (2002). Behavioral research in HIV/AIDS primary and secondary prevention: Recent advances and future directions. *Journal of Consulting and Clinical Psychology*, 70, 626–639.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271–304.
- Kubovy, M., & Healy, A. (1980). Process models of probabilistic categorization. In T. S. Wallsten (Ed.), *Cognitive processes in choice and decision behavior* (pp. 239–262). Hillsdale, NJ: Erlbaum.
- Ladouceur, R., Sylvain, C., Letarte, H., Giroux, I., & Jacques, C. (1998). Cognitive treatment of pathological gamblers. *Behaviour Research and Therapy*, 36, 1111–1120.
- Lee, W. (1971). *Decision theory and human behavior*. New York: Wiley.
- Leigh, B. C. (1989). In search of the seven dwarves: Issues of measurement and meaning in alcohol expectancy research. *Psychological Bulletin*, 105, 361–373.
- Leigh, B. C. (1999). Peril, chance, and adventure: Concepts of risk, alcohol use and risky behavior in young adults. *Addiction*, 94, 371–383.
- Lejuez, C. W., Aklin, W. M., Jones, H. A., Richards, J. R., Strong, D. R., Kahler, C. W., & Read, J. P. (2003). The Balloon Analogue Risk Task (BART) differentiates smokers from nonsmokers. *Experimental and Clinical Psychopharmacology*, 11, 26–33.
- Lejuez, C. W., Aklin, W. M., Zvolensky, M. J., & Pedulla, C. M. (2003). Evaluation of the Balloon Analogue Risk Task (BART) as a predictor of adolescent real-world risk-taking behaviors. *Journal of Adolescence*, 26, 475–479.
- Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., et al. (2002). Evaluation of a behavioral measure of risk-taking: The Balloon Analogue Risk Task (BART). *Journal of Experimental Psychology: Applied*, 8, 75–84.
- Lejuez, C. W., Simmons, B., Aklin, W. M., Daughters, S. B., & Dvir, S. (2004). Risk-taking propensity and risky sexual behavior of individuals in residential substance use treatment. *Addictive Behaviors*, 29, 1643–1647.
- Link, S. W., & Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika*, 40, 77–105.
- Luce, R. D. (2000). *Utility of gains and losses: Measurement-theoretical and experimental approaches*. Mahwah, NJ: Erlbaum.
- Myung, I. J., & Pitt, M. A. (2002). Mathematical modeling. In H. W. Pashler & J. Wixted (Eds.), *Stevens' handbook of experimental psychology: Vol. 4. Methodology in experimental psychology* (3rd ed., pp. 429–460). New York: Wiley.
- National Institute of Mental Health. (2000). *Translating behavioral science into action: Report of the National Advisory Mental Health Council's Behavioral Science Workgroup* (NIMH Publication No. 00–4699). Washington, D.C.: U.S. Government Printing Office. Retrieved August 26, 2005, from www.nimh.nih.gov/publicat/nimhtranslating.pdf
- Nelder, J. A., & Mead, R. (1965). A simplex-method for function minimization. *Computer Journal*, 7, 308–313.
- Onken, L. S., & Bootzin, R. R. (1998). Behavioral therapy development and psychological science: If a tree falls in the forest and no one hears it. *Behavior Therapy*, 29, 539–544.
- Pitt, M. A., Kim, W., & Myung, I. J. (2003). Flexibility versus generalizability in model selection. *Psychonomic Bulletin & Review*, 10, 29–43.
- Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, 109, 472–491.
- Pleskac, T. J. (2004). *Evaluating cognitive sequential risk-taking models: Manipulations of the stochastic process*. Unpublished doctoral dissertation, University of Maryland, College Park.
- Pouget, A., Dayan, P., & Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience*, 26, 381–410.
- Rapoport, A., & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5, 105–120.
- Regenwetter, M., & Marley, A. A. J. (2001). Random relations, random utilities, and random functions. *Journal of Mathematical Psychology*, 45, 864–912.
- Riefer, D. M., Knapp, B. R., Batchelder, W. H., Bamber, D., & Manifold, V. (2002). Cognitive psychometrics: Assessing storage and retrieval deficits in special populations with multinomial processing tree models. *Psychological Assessment*, 14, 184–201.
- Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107, 358–367.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108, 370–392.
- Saad, G., & Russo, J. E. (1996). Stopping criteria in sequential choice. *Organizational Behavior and Human Decision Processes*, 67, 258–270.
- Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the “secretary problem.” *Organizational Behavior and Human Decision Processes*, 69, 221–236.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—Retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4, 145–166.
- Slovic, P. (1966). Risk-taking in children: Age and sex differences. *Child Development*, 37, 169–176.
- Stout, J. C., Busemeyer, J. R., Lin, A. L., Grant, S. J., & Bonson, K. R. (2004). Cognitive modeling analysis of decision-making processes in cocaine abusers. *Psychonomic Bulletin & Review*, 11, 742–747.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect-theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.
- Vickers, D. (1979). *Decision processes in visual perception*. New York: Academic Press.
- Wakker, P., & Stigglebout, A. (1995). Explaining distortions in utility elicitation through the rank-dependent model for risky choices. *Medical Decision Making*, 15, 180–186.
- Weber, E. U., Blais, A.-R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, 15, 263–290.
- Zvolensky, M. J., Lejuez, C. W., Stuart, G. L., & Curtin, J. J. (2001). Experimental psychopathology in psychological science: Historical contributions, integrative features, and future directions. *Review of General Psychology*, 5, 371–381.

Appendix A

The Target Model

We assume that the decision maker (DM) initially selects a target, t_1^* , prior to the first balloon and probabilistically pumps toward it. The DM then adjusts the target down if the balloon explodes or up if it does not prior to his or her stopping, with the magnitude of the adjustments decreasing as the DM gains experience. Specifically, the rule determining the probability that the DM will pump on opportunity i of Balloon h , $r_{h,i}$, given target t_h^* is

$$r_{h,i} = \frac{1}{1 + \exp(\beta \delta_{h,i})}, \quad (A1)$$

where $\delta_{h,i} = i - t_h^*$ and $\beta > 0$. Note that according to this rule, the pumping probability begins high, decreases to .5 on the target opportunity, and continues to decrease. Targets are adjusted on each trial according to the rule

$$t_{h+1}^* = \begin{cases} \text{rnd}[t_h^* + a_1 \exp(-\alpha h)], & \text{if DM stops} \\ \text{rnd}\{(t_h^*[1 - a_2 \exp(-\alpha h)])\}, & \text{if balloon explodes} \end{cases}, \quad (A2)$$

with $a_1, a_2, \alpha > 0$. Thus, the target model has five free parameters, t_1^* , β , a_1 , a_2 , and α . Note the similarity of this model to the hill-climbing model of Busemeyer and Myung (1992).

Appendix B

Constant-Probability Submodel

We assume that the DM optimally updates his or her opinion with experience (i.e., uses Bayes's Rule). The beta-distributed prior opinion is updated following each balloon by successively incrementing m_0 by the number of pumps (i.e., of observations), m_h , made on Balloon h and by successively incrementing a_0 by the number of those pumps, a_h , for which the balloon did not explode. Thus, if the DM stops after a_h pumps on Balloon h and it did not explode, then $m_h = a_h$. If, however, the balloon explodes on Pump m_h , then $a_h = m_h - 1$. The revised estimate of q following experience with Balloon 1 is

$$q_2 = \frac{a_0 + m_1}{m_0 + m_1}$$

if the DM stopped and

$$q_2 = \frac{a_0 + m_1 - 1}{m_0 + m_1}$$

if the balloon exploded. In general, the expression for the DM's estimate of q following experience with h balloons can be written as

$$q_h = \frac{a_0 + \sum_{h'=1}^{h-1} (m_{h'} - d_{h'})}{m_0 + \sum_{h'=1}^{h-1} m_{h'}}, \quad (B1)$$

where $d_{h'} = \begin{cases} 1 & \text{if balloon } h' \text{ exploded} \\ 0 & \text{if balloon } h' \text{ did not explode} \end{cases}$.

And, of course, $p_h = 1 - q_h$.

Appendix C

Increasing-Probability Submodel

To apply Bayes's Rule in this model, we first must develop $p(m_1, d_1, \dots, m_h, d_h|n)$.^{C1} Consider first the case of a sequence of a pumps followed by the balloon's exploding on the last pump. Recalling that

$$p_{hi} = 1/(n - i + 1)$$

and therefore that

$$q_{hi} = 1 - p_{hi} = (n - i)/(n - i + 1),$$

the probability of this sequence given n is

$p(m \text{ pumps ending with an explosion}|n)$

$$= \frac{n-1}{n} \cdot \frac{n-2}{n-1} \cdots \frac{n-m+1}{n-m+2} \cdot \frac{1}{n-m+1} = \frac{1}{n}.$$

Thus, any sequence of pumps resulting in the balloon's exploding has probability $1/n$.^{C2} Similarly, the probability of a sequence of m pumps without the balloon's exploding is

$p(m \text{ pumps ending without an explosion}|n)$

$$= \frac{n-1}{n} \cdot \frac{n-2}{n-1} \cdots \frac{n-m+1}{n-m+2} \cdot \frac{n-m}{n-m+1} = \frac{n-m}{n},$$

subject to the restriction that $n \geq m$. By definition, the above equation equals 0 when $n < m$ because that condition can never occur.

^{C1} We thank Todd Troyer for help with this derivation.

^{C2} Another way to see the same point is to note that the unconditional probability of the balloon's popping at any pump is $1/n$.

Thus, in general the result can be expressed as

$$p(m_1, d_1, m_2, \dots, m_h, d_h) = \prod_{h'=1}^h \left[\frac{s_{h'}(n - m_{h'})^{(1-d_{h'})}}{n^{h'}} \right],$$

where $d_{h'}$ is defined as it was for Equation B1 and

$$s_{h'} = \begin{cases} 0 & \text{if } n < m^* \\ 1 & \text{if } n \geq m^* \end{cases}$$

with $m^* = \max(m_1, m_2, \dots, m_h)$.

We can now obtain $p(n|m_1, d_1, \dots, m_h, d_h)$ with Bayes's rule. The result is

$$p(n|m_1, d_1, \dots, m_h, d_h) = \frac{\prod_{h'=1}^h [s_{h'}(n - m_{h'})^{(1-d_{h'})} p(n)/n^{h'}]}{\sum_{n'=1}^{\infty} \prod_{h'=1}^h [s_{h'}(n' - m_{h'})^{(1-d_{h'})} p(n')/n'^{h'}]}. \quad (C1)$$

From this equation, it is straightforward to obtain the expected value of the posterior distribution following Balloon h as

$$\hat{n}_h = \sum_{n'=1}^{\infty} n' p(n'|m_1, d_1, \dots, m_h, d_h). \quad (C2)$$

This result, then, leads to the DM's SP of Balloon h exploding on opportunity i ,

$$p_{hi} = \frac{1}{\hat{n}_h - i + 1}. \quad (C3)$$

Appendix D

Maximum Likelihood Estimation (MLE)

Equation 2 provides the likelihood equation for the baseline model, but it does not serve as a general likelihood equation for all models. Instead, we write the product of the response likelihoods over all pump opportunities and balloons (assuming independence) as

$$L(a_1, d_1, a_2, d_2, \dots, a_h, d_h, \dots, a_f, d_f)$$

$$= \prod_{h=1}^f \prod_{i=1}^{c_h} \hat{r}_{h,i} (1 - \hat{r}_{h,c_h+1})^{d_h}, \quad (D1)$$

where f is the total number of balloons ($f = 30$ in the present data set), c_h is the total number of pumps the DM took on Balloon h , and $d_h = 1$ if the DM stopped on opportunity $c_h + 1$, 0 otherwise.

Equation D1 is optimized separately for each DM for each model. This is done by first taking the log of the equation (the log likelihood, primarily for computational ease, but also in many cases for statistical purposes), setting $\hat{r}_{h,i} = r_{h,i}$, substituting the model equations for $r_{h,i}$ and finally searching for the parameter values within the model that maximize the equation for each DM. Although this is a straightforward matter for the baseline model, estimating the log likelihoods for the nine other theoretical models is not easy, as none has an analytical solution yielding a global maximum. Thus, we used computer optimization routines to estimate the maximum likelihood solutions for each model per DM. Numerous routines are available, with each performing best under different conditions. After testing several, including a sequential quadratic programming routine and variants of genetic algorithms, we found that the Nelder–Mead numerical optimization routine (Nelder & Mead, 1965), available in Matlab, generally

performed the best. However, as with all optimization routines, the success of the Nelder–Mead method depends on the vector of starting values when the to-be-optimized function has local maxima, as ours do because of their highly irregular surfaces.

To combat this problem, we used a grid-search technique to select numerous starting values as input to the Nelder–Mead method. Specifically, for each model, we first selected a plausible space for each parameter and divided it into thirds.^{D1} (For example, for γ^+ from the evaluation models, we considered the space between 0 and 1.35 to be plausible; each interval then had a width of 0.45.) The computer then randomly chose one of the hypercubes defined by the boundaries of the parameter intervals and randomly selected a parameter vector from it. The vector was tested to ensure that its values would lead to a solution above a prespecified criterion of $1nL > -2,000$. If it did, then it was fed to the Nelder–Mead method as the starting vector. If not, then another hypercube was randomly chosen and the procedure was begun again. We used 50–100 iterations of this full procedure to obtain the maximum log-likelihood estimate for each model per DM.

^{D1} This space was only used to select starting values. The Nelder–Mead method was allowed to search among the full parameter space defined for each model.

Received August 14, 2004
Revision received May 3, 2005
Accepted May 9, 2005 ■