



Observing evidence accumulation during multi-alternative decisions

Scott Brown^{a,*}, Mark Steyvers^b, Eric-Jan Wagenmakers^c

^a School of Psychology, University of Newcastle, Callaghan NSW 2308, Australia

^b Department of Cognitive Science, University of California, Irvine, United States

^c Department of Psychology, University of Amsterdam, The Netherlands

ARTICLE INFO

Article history:

Received 10 June 2009

Received in revised form

20 August 2009

Available online 13 October 2009

Keywords:

Probabilistic inference

Hick's Law

Optimal observer models

Speed–accuracy tradeoff

ABSTRACT

Most decision-making research has focused on choices between two alternatives. For choices between many alternatives, the primary result is Hick's Law—that mean response time increases logarithmically with the number of alternatives. Various models for this result exist within specific paradigms, and there are some more general theoretical results, but none of those have been tested stringently against data. We present an experimental paradigm that supports detailed examination of multi-choice data, and analyze predictions from a Bayesian ideal observer model for this paradigm. Data from the experiment deviate from the predictions of the Bayesian model in interesting ways. A simple heuristic model based on evidence accumulation provides a good account for the data, and has attractive properties as a limit case of the Bayesian model.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Models of human decision-making are both prolific and varied. A common thread arises, however, when researchers consider not just the choices that people make, but also the time it takes people to make these choices. Almost without exception, decision-making models that include predictions about choice latency (i.e., response time, RT) employ the notion of “evidence accumulation” (e.g. Brown & Heathcote, 2008; Diederich & Busemeyer, 2006; Luce, 1986; Ratcliff, 1978; Ratcliff & McKoon, 2008; Ratcliff & Smith, 2004; Roe, Busemeyer, & Townsend, 2001; Townsend & Ashby, 1983; Van Zandt, Colonius, & Proctor, 2000; Vickers & Lee, 1998; Wagenmakers, van der Maas, & Grasman, 2007). In evidence accumulation models, a decision is made by repeatedly sampling information from the environment and considering this as evidence for or against the different choice alternatives. This stream of information is accumulated until sufficient evidence is gathered to trigger a decision. In experimental practice, the decisions being made are typically very fast (often measured in milliseconds), and the environmental stimulus is almost always held constant. This means that the stream of information is internal to the decision-maker, and cannot be directly observed or measured. Some efforts have been made to inspect the stream of evidence more closely, for example by querying the amount of accumulated evidence at

various times (Meyer, Irwin, Osman, & Kounios, 1988). It is not entirely clear, however, whether such methods actually allow the inspection of internal evidence (Ratcliff, 1988). Other researchers have investigated the theoretical consequences of moment-by-moment changes in the way evidence is accumulated. Such theories have provided successful accounts of choices in which many different stimulus attributes must be considered, perhaps sequentially (Diederich, 1997), and perceptual choices in which attention is manipulated by masking (Smith & Ratcliff, 2009).

There have been several attempts to go further, and directly manipulate the stream of evidence accumulated by the decision-maker, using an “expanded judgment” paradigm developed by Irwin, Smith, and Mayfield (1956). Brown and Heathcote (2005), Pietsch and Vickers (1997) and Usher and McClelland (2001) studied two-choice decision processes using expanded judgments, in which the information provided to the observer was manipulated moment by moment. In each case, the two response choices corresponded to two different perceptual properties of the stimulus: horizontal location, for Pietsch and Vickers; letter identity for Usher and McClelland; and orientation for Brown and Heathcote. Their experiments used stimuli that rapidly switched between favoring one response choice or the other, and the decision-maker's task was to evaluate the overall weight of evidence.

All three experiments were used for the same purpose—to examine the role of memory in accumulating evidence. Memory plays an important part in this task because the observer is presented with a constantly varying stream of information, and evidence accumulation requires them to recall elements of that

* Corresponding author.

E-mail address: scott.brown@newcastle.edu.au (S. Brown).

stream presented earlier. For example, Pietsch and Vickers found that this memory load imposed a severe limit on evidence accumulation: information provided early in the decision process was forgotten later. In terms of the expanded judgment task, what Pietsch and Vickers found was that decisions were more heavily influenced by information provided later in the stimulus stream rather than earlier. Brown and Heathcote found that this “leakiness” of evidence accumulation was decreased by extensive practice, and Usher and McClelland identified large individual differences in the dynamics of evidence accumulation.

We extend these investigations of the evidence accumulation process by proposing a new paradigm that retains the essential features of the previous studies, but adds some key improvements. In particular, our paradigm completely lifts the memory requirements imposed on the observer, transferring them to the stimulus display instead: the entire history of the evidence stream is always available on screen. This allows us to investigate elements of the evidence accumulation process (such as the rules employed for terminating the accumulation process) independent of memory limitations. Our paradigm also naturally allows for participants to continue collecting information samples as long as they please, and make a decision only when they feel confident. This “information-controlled” paradigm simultaneously generates data on both response accuracy and decision time, providing greater constraint than prior investigations, which focused on either fixed decision times, fixed (high) accuracy, or a fixed number of evidence samples.

We also set out the ideal observer model corresponding to our paradigm, and we provide an example application that involves Hick’s Law (Hick, 1952). Finally, we discuss a sub-optimal heuristic model that provides a good approximation to the empirical data, and which also has an interesting relationship to the ideal observer model.

2. The paradigm

At the start of every trial of the experiment, the participant is shown a display that features K empty columns (see Fig. 1 for an illustration). During the trial, time proceeds in discrete steps of equal duration. Throughout, when we refer to “decision time”, the measurement unit is this step value (0.2 s in our experiment). Small objects (henceforth “bricks”) fall from the top of the screen and accumulate on the columns. During each time step, each column may accumulate one extra brick, or none. One of the columns has a higher accumulation rate than the others, and the participant’s goal is to select that target column as quickly as possible without making too many mistakes.¹

Many properties of this paradigm can be manipulated to create informative experiments. We focus here on manipulating the number of choice alternatives (the number of columns, K). We also manipulate response urgency, by sometimes instructing participants to be cautious, and sometimes to be fast (i.e., the speed–accuracy tradeoff, e.g., Forstmann et al., 2008; Hick, 1952; Schouten & Bekker, 1967).

2.1. Notation

One column accumulates bricks faster than the other columns; we call this the *target* column and the others *distractor* columns. At

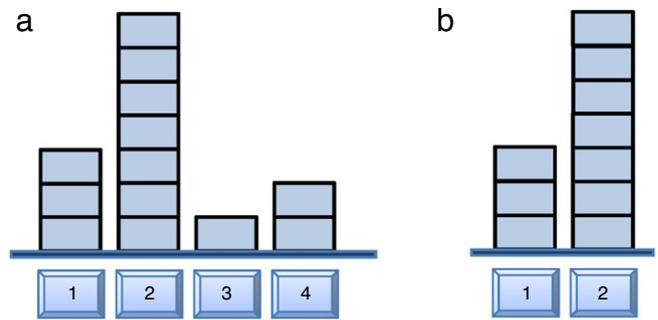


Fig. 1. Screenshot of the experimental paradigm after 10 time points have passed. Panel (a) depicts an example four-choice situation and panel (b) depicts an example two-choice situation. The goal for the participants is to identify the one column that has a higher accumulation rate than the others. At each time point, participants can either wait and observe more information, or they can make a choice by pressing one of the buttons below the columns. See text for details.

each time step, a column might or might not accumulate a brick. The probability of accumulating an extra brick at any time step is independent across columns and time steps. For the target column, the probability of accumulating a brick at any time step is denoted by $\theta_{(t)}$, and for all of the distractor columns it is denoted by $\theta_{(d)}$. Note that we assume an equal accumulation probability for all distractor columns. When discussing decision models, we denote the hypothesis that the i th column is the target by H_i . We use D to denote the observed data. For a particular column i , s_i is the number of “successes” (i.e., accumulated bricks) and $f_i = n - s_i$ is the number of “failures” (i.e., missed bricks), at time step n .

3. Multi-alternative choices and Hick’s law

Our experimental paradigm is naturally suited to study the decisions between any number of alternatives, from $K = 2$ up to however many columns will neatly fit on the display screen (we stopped at $K = 12$). In contrast, most other decision-making research has focused on binary choices ($K = 2$); even in that apparently simple paradigm, the richness of empirical results is surprising. For choices between more than two alternatives, the fundamental result is Hick’s Law (Hick, 1952; Hyman, 1953; Merkel, 1885; for a review see Teichner & Krebs, 1974). In its simplest form, Hick’s Law states that the mean RT is linear in the *logarithm* of the number of choice alternatives K : mean RT = $a + b \log(K)$, where a and b are constants.²

Hick motivated his results by recourse to information theory (Shannon & Weaver, 1949), arguing that “the amount of information extracted is proportional to the time taken to extract it, on the average” (1952, p. 25); for alternative accounts and critique see Christie and Luce (1956) and Laming (1966, 1968). Although Hick discussed a binary search algorithm that would be consistent with the data, he concluded that “At present... it is impossible to venture beyond the general statement in terms of information theory. This, indeed, may be adequate for practical applications; but it inevitably leaves the details vague; and so they must remain, until more evidence or better reasoning is brought to the problem.” (Hick, 1952, p. 25).

In recent years, the problem of choosing between multiple alternatives has been considered from several different theoretical perspectives. Brown and Heathcote (2008), Brown, Marley, Donkin, and Heathcote (2008), Roe et al. (2001), Usher and McClelland

¹ An online version of the experiment can be found at <http://psiexp.ss.uci.edu/~yoshi/expW2.html>.

² Some authors add a constant value to K , but we have not found that necessary in our analyses.

(2001), and Vickers and Lee (2000) have developed evidence accumulation models for multi-alternative choice for particular paradigms (i.e., consumer decisions, perceptual choice, or absolute identification). Even though they are focused on paradigm-specific outcomes, each of these models has been shown to be consistent with Hick’s Law, as have data from the associated paradigms. Usher, Olami, and McClelland (2002) (see also Bogacz, Usher, Zhang, & McClelland, 2007 and McMillen & Holmes, 2006) took a more general approach to evidence accumulation and Hick’s Law. They showed that a standard accumulator model was able to predict Hick’s Law if the decision threshold (i.e., the amount of evidence required to initiate a decision) was allowed to increase as the logarithm of the number of choice alternatives. This assumption is reasonable if participants strive to maintain a constant error rate in the face of increasing response alternatives (Usher et al., did not discuss whether participants actually do so strive).

Several authors have considered the problem of multi-alternative decision-making from a purely statistical point of view (e.g., Baum & Veeravalli, 1994; Dragalin, Tartakovsky, & Veeravalli, 1999, 2000). In particular, Baum and Veeravalli (1994) developed the multi-hypothesis sequential probability ratio test (MSPRT), a statistical procedure that generalizes the famous sequential probability ratio test or SPRT (Wald & Wolfowitz, 1948) to more than two-choice alternatives.³ This work was then extended by Dragalin et al. (1999, 2000), who also demonstrated that a simple heuristic approach performs just as well as the statistically optimal algorithm, at least when error rates are low. Below, we develop particular versions of these ideal observer models, appropriate to our decision paradigm. We then describe some empirical results and later discuss the performance of a heuristic model in describing the data.

4. Two optimal observer models

A truly optimal analysis of our paradigm would take into account the relative costs of making errors (choosing the wrong column) vs. the cost of waiting longer, and sampling more information (Baum & Veeravalli, 1994; Berger, 1985; Berry & Fristedt, 1985). To simplify matters, we instead content ourselves with a restricted definition: a model is “optimal” if, for some predetermined error rate, the expected decision time is minimized. In this sense, an optimal decision-maker faces the following question on each time step: “do I want to stop sampling and select the highest column (i.e., H_h) or do I want to observe more information?” The ideal observer, therefore, computes $p(M_h|D)$, and decides to choose column h when this probability exceeds a criterion number. This is the MSPRT procedure that generalizes the SPRT procedure to more than two-choice alternatives.

According to Bayes’ theorem,

$$p(H_h|D) = \frac{p(D|H_h)p(H_h)}{\sum_j p(D|H_j)p(H_j)}. \tag{1}$$

When the *a priori* probabilities for the columns are equal—an assumption that we make from now on—this simplifies to

$$p(H_h|D) = \frac{p(D|H_h)}{\sum_j p(D|H_j)}. \tag{2}$$

³ Bogacz and Gurney (2007) showed how the MSPRT can be implemented in the brain.

4.1. Model 1: Probabilities-known-exactly

In the first analysis, assume that $\theta_{(t)}$ and $\theta_{(d)}$ are known exactly. (This might be close to true for participants after some practice.) Consider first H_h . Under this hypothesis, the data s_h from n samples originated from a binomial process with parameter $\theta_{(t)} \in [0, 1]$: $p(s_h, n|H_h) = \binom{n}{s_h} \theta_{(t)}^{s_h} (1 - \theta_{(t)})^{f_h}$, where $f_h = n - s_h$. The data for each of the remaining columns, s_j with $j \neq h$, originated from a binomial process with parameter $\theta_{(d)}$. Together, this means that $p(D|H_h) = \binom{n}{s_h} \theta_{(t)}^{s_h} (1 - \theta_{(t)})^{f_h} \prod_{j \neq h} \binom{n}{s_j} \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j}$. Similar calculations hold for the hypotheses that any other column is the one with the highest rate, yielding

$$\begin{aligned} p(H_h|D) &= \frac{\binom{n}{s_h} \theta_{(t)}^{s_h} (1 - \theta_{(t)})^{f_h} \prod_{j \neq h} \binom{n}{s_j} \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j}}{\sum_k \left\{ \binom{n}{s_k} \theta_{(t)}^{s_k} (1 - \theta_{(t)})^{f_k} \prod_{j \neq k} \binom{n}{s_j} \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j} \right\}} \\ &= \frac{\theta_{(t)}^{s_h} (1 - \theta_{(t)})^{f_h} \prod_{j \neq h} \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j}}{\sum_k \left\{ \theta_{(t)}^{s_k} (1 - \theta_{(t)})^{f_k} \prod_{j \neq k} \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j} \right\}} \\ &= \frac{\theta_{(t)}^{s_h} (1 - \theta_{(t)})^{f_h} \left[\prod_j \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j} \right]}{\theta_{(d)}^{s_h} (1 - \theta_{(d)})^{f_h}} \\ &= \frac{\sum_k \left\{ \theta_{(t)}^{s_k} (1 - \theta_{(t)})^{f_k} \left[\prod_j \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j} \right] \right\}}{\sum_k \left\{ \theta_{(d)}^{s_k} (1 - \theta_{(d)})^{f_k} \left[\prod_j \theta_{(d)}^{s_j} (1 - \theta_{(d)})^{f_j} \right] \right\}} \\ &= \frac{\left[\theta_{(t)}^{s_h} (1 - \theta_{(t)})^{f_h} \right]}{\left[\theta_{(d)}^{s_h} (1 - \theta_{(d)})^{f_h} \right]} \\ &= \frac{\sum_k \left\{ \left[\theta_{(t)}^{s_k} (1 - \theta_{(t)})^{f_k} \right] \right\}}{\sum_k \left\{ \left[\theta_{(d)}^{s_k} (1 - \theta_{(d)})^{f_k} \right] \right\}} \\ &= \frac{(\theta_{(t)}/\theta_{(d)})^{s_h} \left[(1 - \theta_{(t)}) / (1 - \theta_{(d)}) \right]^{f_h}}{\sum_k \left\{ (\theta_{(t)}/\theta_{(d)})^{s_k} \left[(1 - \theta_{(t)}) / (1 - \theta_{(d)}) \right]^{f_k} \right\}}, \tag{3} \end{aligned}$$

where k indexes the hypothesis entertained (i.e., the hypothesis that column k is the target column), ranging from 1 to K , the total number of columns.

The optimal observer would monitor $p(H_h|D)$ and initiate a response when a criterion level c is exceeded.

4.1.1. Example

As an example of the decisions made by the optimal observer, consider the two experimental situations shown in Fig. 1. In both cases, the first two columns have the same number of bricks. The difference is that the second example has two additional columns. The optimal observer makes different predictions for these two cases—the additional columns add to the uncertainty of which column corresponds to the target column. Suppose we use the binomial rate parameters $\theta_{(t)} = .50$ and $\theta_{(d)} = .35$. As reported later, these rates correspond to the rates used in the experiment with human participants. Eq. (3) leads to $p(H_h|D) = [0.078 \ 0.923]$ and $p(H_h|D) = [0.073 \ 0.867 \ 0.021 \ 0.039]$ for the two cases, respectively. Therefore, if a criterion setting $c = 0.9$, is used (see below), the first case would trigger a response (picking the second column) but the second case would not.

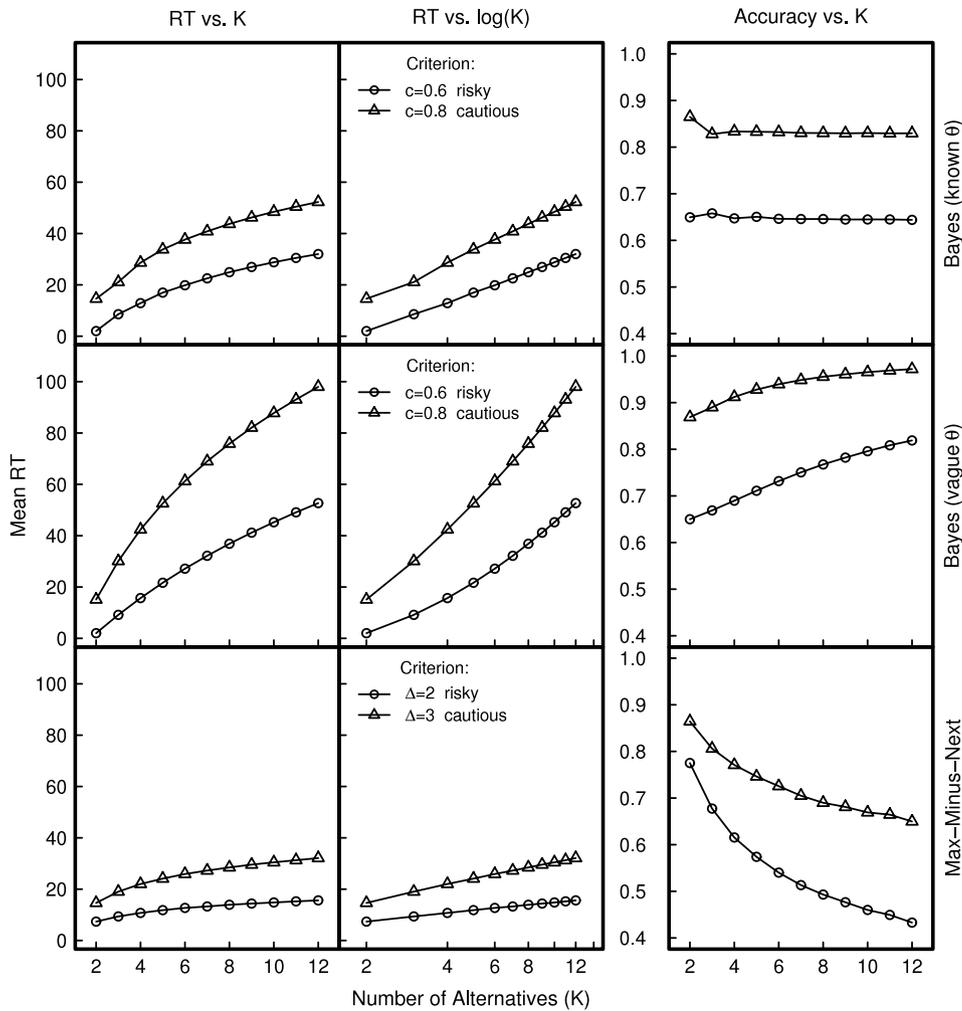


Fig. 2. Mean RT and accuracy simulated from the three models. The top row illustrates the “probabilities-known-exactly” model; the middle row illustrates the “probabilities-known-vaguely” model; and the bottom row illustrates the heuristic “max-minus-next” model, discussed later. The left panels show mean decision time (measured in discrete time steps) as a function of the number of choice alternatives, for two different criterion settings (risky or cautious). The center panels show the same data with a logarithmic abscissa. The right panels plot mean choice accuracy against the number of choice alternatives, K .

4.1.2. Simulations

To systematically explore the behavior of the model, we varied the number of choice alternatives and criterion settings and simulated the task a large number of times. The data in each trial were generated probabilistically according to the binomial rate parameters $\theta_{(t)} = .50$ and $\theta_{(d)} = .35$. The number of choice alternatives K ranged from 2 to 12 (i.e., $K \in \{2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$ —note that our participants experienced only $K \in \{2, 4, 8, 12\}$).

The criterion parameter, c , sets a response threshold on the posterior probability $p(H_h|D)$, and thus represents the accuracy that the model aims to attain; to the extent that the statistical structure of the environment is faithfully represented by the equations above, the model will perform with overall accuracy c . We suppose, initially, that c is set by the participant at a fixed level (independent of K), and that this level is influenced by metacognition beyond the scope of the present article, such as the desire to perform well, and the competing desire to finish the experiment quickly, and also by the instructions given to the participant. We studied the behavior of the model under two criterion settings: risky ($c = .6$), and cautious ($c = .8$). These criterion settings mirror the instructions that participants received in our experiment, but the absolute magnitude of a “risky” and a “cautious” criterion setting will change from situation to situation, depending on the cost of an incorrect response. For each combination of K and c , the optimal observer made one million decisions (i.e., $11 \times 2 = 22,000,000$ in total). The two

dependent measures of interest were the proportion of correct target identifications and the number of observations needed before committing a decision.

Fig. 2 shows the results of these simulations, for the current model (top row of Fig. 2) and for two models to be discussed later (middle and lower rows). For each model (i.e., in each row) the left panel shows how mean decision time changes with the number of choices (K), and the middle panel shows this same relationship using a logarithmic abscissa. The right panel shows how accuracy changes with the number of choices. The model is well calibrated in that it predicts a response accuracy that is determined by the threshold set on the posterior probability, either $c = .6$ or $c = .8$. This illustrates the model’s optimality—choosing the maximum column any sooner than this model must result in more errors than the criterion amount set by the c parameter. There is a small amount of overshoot error, however—the predicted response accuracy rates are slightly above the calibrated values set by the parameter c . This overshoot is a consequence of the non-zero size of the discrete time step used in our paradigm. The model issues a response at the first time step on which the threshold c is reached or exceeded. The center panel on the top row of Fig. 2 shows that Hick’s Law is predicted by this model, even though it was not designed for that purpose. That is, the predicted mean RT increases log-linearly with the number of decision alternatives. Note that, due to the model’s calibrated nature, predicted response accuracy does not change with the number of decision alternatives.

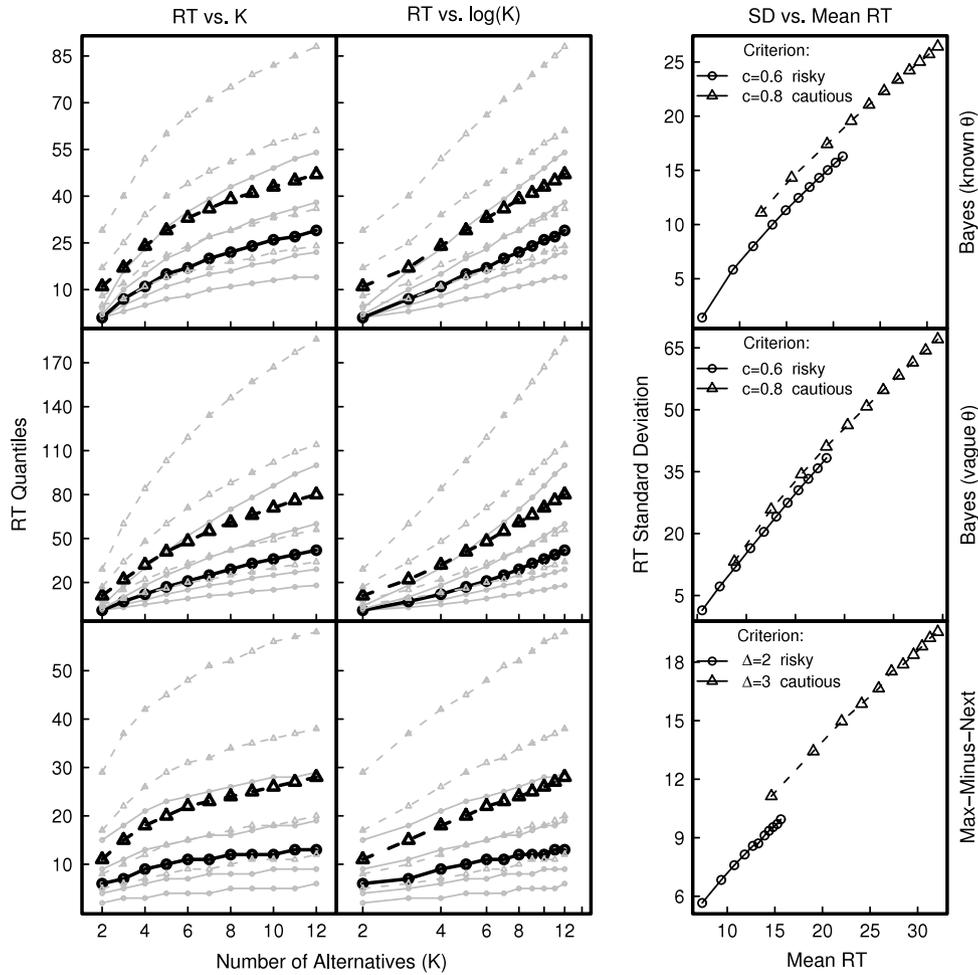


Fig. 3. RT quantiles, means, and standard deviations simulated from the three models: top row shows the “probabilities-known-exactly” Bayesian model; middle row shows the “probabilities-known-vaguely” Bayesian model; and the bottom row shows the heuristic “max-minus-next” model. In each row, the left panel shows five percentiles of the decision time distribution (10th, 30th, 50th, 70th and 90th) as functions of the number of choice alternatives, for two different criterion settings (risky and cautious). The 50th percentile is the median, and is shown as a heavy black line; the others are grey. The center panels show the same data with a logarithmic abscissa. The right panels plot the standard deviation of the RT distributions against their means.

Fig. 3 analyzes Hick’s Law in greater detail, by examining RT distribution features, rather than just the mean. Once again, the three rows correspond to three different models: the top row corresponds to the current Bayesian model, the middle and lower rows correspond to models introduced below. Within each row, the left and center panels illustrate the usual Hick’s Law dependence of RT on the logarithm of the number of choice alternatives, except that this time RT is measured by response quantiles. These quantiles summarize the full RT distributions using five percentiles: 10th, 30th, 50th (i.e., the median, drawn in heavier face), 70th and 90th. Thus, the model generates Hick’s Law not just for the RT mean, but also for the RT quantiles, although there may be some non-linearity in the logarithmic plot for the smallest choice set sizes ($K = 2, 3$). The right hand panel in each row shows the predicted relationship between the standard deviation of the RT distribution and its mean. The current model (top row) predicts a linear relationship between these variables, as often observed empirically (Wagenmakers & Brown, 2007). Interestingly, the linear function relating standard deviation and mean appears to be almost identical for cautious vs. risky responses, under this model.

4.2. Model 2: Probabilities-known-vaguely

In the second analysis, we assume that the participant has only a noisy representation of the underlying rates, described by

a beta distribution with parameters α and β . We also assume that the distractors all have the same (vaguely known) rate. Thus, the knowledge about a fixed $\theta_{(t)}$ and a fixed $\theta_{(d)}$ is now replaced by beta distributions with parameters $(\alpha_{(t)}, \beta_{(t)})$ and $(\alpha_{(d)}, \beta_{(d)})$, respectively (see Fig. 4). Dropping for the moment the subscripts for t and f , this means that

$$p(\theta) = \frac{1}{B(\alpha, \beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1}, \tag{4}$$

where B is the beta function, defined as $B(\alpha, \beta) = \int_0^1 x^{\alpha-1} (1 - x)^{\beta-1} dx$, or as $B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$, where Γ is the gamma function, which for a positive integer n simplifies to $\Gamma(n) = (n - 1)!$.

In order to compute $p(D|H_h)$, say, we need to integrate out the model parameter θ , weighting the probability of observing the data given a particular value of θ with the prior probability of θ as given in Eq. (4). Consider for instance the probability of observing data s_h from n samples. We then have

$$\begin{aligned} p(s_h) &= \int_0^1 p(\theta_{(t)}) p(s_h | \theta_{(t)}) d\theta_{(t)} \\ &= \binom{n}{s_h} \frac{1}{B(\alpha, \beta)} \int_0^1 \theta_{(t)}^{\alpha+s_h-1} (1 - \theta_{(t)})^{\beta+f_h-1} d\theta_{(t)} \\ &= \binom{n}{s_h} \frac{1}{B(\alpha, \beta)} B(\alpha + s_h, \beta + f_h). \end{aligned} \tag{5}$$

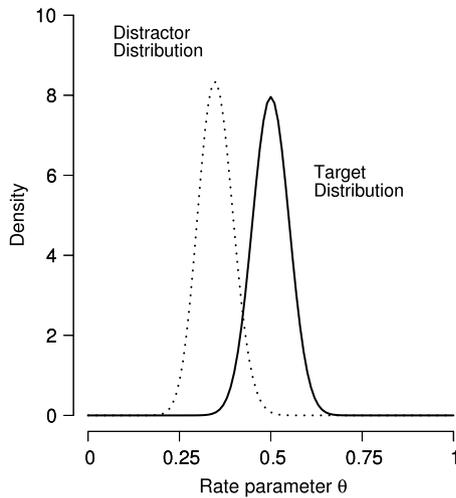


Fig. 4. Distributions that quantify uncertainty for the rate parameters. The simulations reported in the text assume a Beta($\alpha_t = 50, \beta_t = 50$) distribution for the target column (solid line) and a Beta($\alpha_d = 35, \beta_d = 65$) distribution for the distractors (dotted line).

This means that Eq. (3) can be adjusted to account for parameter uncertainty as follows:

$$p(H_h|D) = \frac{B(\alpha_t + s_h, \beta_t + f_h)B(\alpha_d + \sum_{j \neq h} s_j, \beta_d + \sum_{j \neq h} f_j)}{\sum_k \left\{ B(\alpha_t + s_k, \beta_t + f_k)B(\alpha_d + \sum_{j \neq k} s_j, \beta_d + \sum_{j \neq k} f_j) \right\}}. \quad (6)$$

4.2.1. Example

Consider again the two experimental situations shown in Fig. 1. We simulated the model for these two cases, but instead of using a fixed probability of $\theta_t = .50$ and $\theta_d = .35$, we used two beta distributions that quantify uncertainty with respect to these rates. As shown in Fig. 4, we used Beta($\alpha_t = 50, \beta_t = 50$), for the target column and Beta($\alpha_d = 35, \beta_d = 65$) for the distractor columns. Eq. (6) leads to $p(H_h|D) = [0.094 \ 0.906]$ and $p(H_h|D) = [0.052 \ 0.898 \ 0.019 \ 0.031]$ for the two cases, respectively. This outcome is similar to the prediction made by Model 1, but with somewhat decreased confidence about the second column corresponding to the target column. This is due to the uncertainty in the underlying rates used for the target and distractors.

4.2.2. Simulations

We carried out the same simulations as before to explore the effect of the number of choice alternatives and criterion settings.

The middle rows of Figs. 2 and 3 illustrate that, when compared with the model with exact knowledge, the model with vague knowledge needs considerably more samples to attain a comparable level of accuracy. Moreover, the model with vague knowledge only approximates Hick's Law, as the function that relates mean RT to $\log(K)$ is slightly convex instead of straight.

Note that this model fails to predict a calibrated error rate. The right panel of the middle row of Fig. 2 shows that the vague-knowledge Bayesian model predicts increasing accuracy with increasing number of choice alternatives (at the expense of greatly increased decision time). At first sight, this result is surprising, since the model was developed to maintain a constant posterior probability. The apparent problem arises because of model-misspecification: the stimuli on which the simulations were based mirrored those used in our experiment, not those assumed by the

vague-knowledge model—the simulations used fixed values of θ_t and θ_d , not beta-distributed values.

Finally, the right panel in the second row of Fig. 3 shows that, just as the model with exact knowledge, the model with vague knowledge predicts a linear relationship between RT mean and RT standard deviation.

4.3. Conclusions from the optimal observer models

The ideal observer models present some interesting findings. They demonstrate that Hick's Law can be produced naturally by a model using a fixed evidence threshold across different numbers of response alternatives. The models also make some interesting predictions for data, for example that Hick's Law should hold not just for mean RT, but also for each quantile of the RT distribution. The models also predict a linear relationship between the mean and the standard deviation of the RT distribution (as observed by Wagenmakers & Brown, 2007, for other paradigms). This prediction is at odds with prior work by both Hick (1952) and Laming (1966), which tentatively suggests that RT mean increases linearly with RT variance and not with RT standard deviation.

5. Experiment

5.1. Methods

Thirty-nine undergraduates from the University of California, Irvine, completed an experiment using the paradigm described above. There were 32 trials in each of eight blocks. Participants were instructed to make very careful (but slower) decisions during the first four blocks, and faster (but less accurate) decisions during the last four blocks (or vice versa, counterbalanced across participants). On each trial, an array of columns was set out across the bottom of the screen, representing the different response alternatives. The number of response alternatives was chosen randomly on each trial from $K \in \{2, 4, 8, 12\}$, subject to each value of K occurring equally often in every block. The columns were initially empty—of height zero. On each discrete time step (200 ms) a block fell on each distractor column with probability $\theta_d = .35$ and on the target column with probability $\theta_t = .5$ (the target column was chosen randomly on each trial). The participant's task was to choose the target column, by clicking on a button below it, as illustrated in Fig. 1. Feedback was given after each trial. Instructions emphasized the random nature of the process, and that a column which appeared to be the target column early in a trial may later turn out to be a distractor. The instructions also suggested that all distractor columns were statistically equal.

An ideal version of this paradigm would allow infinite evidence accumulation—participants could always make perfectly accurate decisions by waiting long enough. A practical problem is that, after some number of blocks have accumulated, column heights grow near the top of the display. We solved this problem by having the column display window smoothly re-size by reducing the scale whenever a column grew too large. Other solutions are possible for this problem, but we settled on this version because it kept the critical feature that participants could always wait longer, and collect more evidence, if they chose.

5.2. Results

Fig. 5 illustrates the data using the same format as used for each model (i.e., each row) of Fig. 2. The error bars represent ± 1 standard error, calculated using the methods of Loftus and Masson (1994). Similarly, Fig. 6 illustrates the distribution details observed in the data, using the same format as used for each row of Fig. 3. By using the same graph format throughout, the various model

Fig. 5. Mean RT and response accuracy from the experiment, using the same format as Fig. 2. The grey curve in the right panel shows chance performance level.

Fig. 6. Observed RT quantiles and the relationship between standard deviation and mean RT, using the same format as Fig. 3.

Table 1

Mean percentage (and SD) of times that participants chose the column with the greatest instantaneous height, separately for each number of choice alternatives (K) and for speed-emphasis and accuracy-emphasis instructions.

	$K = 2$	$K = 4$	$K = 8$	$K = 12$
Speed	96(3.4)	90(6.1)	85(9.8)	75(11)
Accuracy	97(4.0)	97(4.5)	92(9.1)	87(11)

predictions can easily be compared with the data by comparing Fig. 5 against each row of Fig. 2 and comparing Fig. 6 against each row of Fig. 3.

The left and center panels of Fig. 5 confirm that we observed Hick's Law in its usual format—mean RT increased linearly with the logarithm of the number of choice alternatives. This relationship held both under speed-emphasis and accuracy-emphasis instructions. The findings in the right panel of Fig. 5 are quite different from classical investigations of Hick's Law; they show that the response accuracy decreased markedly as the number of alternatives increased, although performance was always well above the chance level (the grey line). In the earlier examinations of Hick's Law using information-controlled decisions (such as Experiment 1 of Hick, 1952), experimental measures were taken to ensure accuracy remained perfect, at all set sizes. These measures allowed researchers to consider only one dependent variable (RT), rather than the joint effects of RT and accuracy. In more recent studies, both accuracy and RT have been free to vary—as in our study—and accuracy has been observed to decrease with larger numbers of choice alternatives, at least with keypress responses (e.g., Kveraga, Boucher, & Hughes, 2002; Lee, Keller, & Heinen, 2005). Our theoretical models unify these dependent variables, simultaneously explaining changes in both.

The RT distributions, illustrated by quantile estimates in Fig. 6 suggest that the log-linear relationship of Hick's Law also hold for the entire RT distribution, not just its mean. This is consistent

with predictions from the optimal observer models, and also with data from the saccadic-response choice task used by Lee et al. (2005). The right panel of Fig. 6 shows also that there is at least an approximately linear relationship between the mean of RT and its standard deviation. The empirical data also include an element missing from optimal models—participants did not always choose the column with the largest height, see Table 1. This is, of course, sub-optimal, because the largest column is the one most likely to have the higher accumulation rate. It is possible that these errors were caused either by perceptual limitations (it may be difficult to notice which column is larger in a display of many columns) or by a time lag between decision and response execution (a decision is made, selecting the tallest column, but the second-tallest column grows larger in the moments before the response button press is executed). Our data tentatively support the time-lag hypothesis, by considering the state of the display a few time steps *before* each response. The chosen column was the tallest column at the time of responding on only 89.8% of trials. However, the chosen column was tallest on almost every trial (97.2%), if one looks back to three time steps before the response. This suggests a delay between choosing to respond, and actually making the response. The proportion of trials for which the chosen column was tallest prior to the response lag (i.e., 97.2%) was larger for the estimated lag (three time steps, or 0.6 s) than for any other delay value. We have also re-analyzed the data (not shown in the figures), after removing all trials where participants did not pick the tallest column. This analysis did not change any of the qualitative patterns we observe for the accuracy results in Fig. 5. Therefore, whatever causes participants to choose an apparently sub-optimal column does not appear to change the important behavioral patterns.

6. A heuristic model: max-minus-next

Dragalin et al. (1999, 2000) showed that a very simple decision heuristic was able to perform just as well as the optimal MSPRT

