

Robust versus optimal strategies for determining the speed-accuracy tradeoff on two-alternative forced choice tasks

M. Zacksenhouse¹, P. Holmes² and R. Bogacz³.

¹ Faculty of Mechanical Engineering, Technion – Israel Institute of Technology, Haifa 32000, Israel.

² Program in Applied and Computational Mathematics and Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, U.S.A.

³ Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.

July 9, 2008

Abstract

It has been proposed that animals and humans might choose a speed-accuracy tradeoff that maximizes reward rate. For this utility function the simple drift-diffusion model of two-alternative forced-choice tasks predicts a parameter-free optimal performance curve that relates normalized decision times to error rates. However, behavioral data indicate that only $\approx 30\%$ of subjects achieve optimality, and here we investigate the possibility that, in allowing for uncertainties, subjects might exercise robust strategies instead of optimal ones. We consider two strategies in which robustness is achieved by relinquishing performance: maximin and robust-satisficing. The former supposes maximization of guaranteed performance under a presumed level of uncertainty; the latter assumes that subjects require a critical performance level and maximize the level of uncertainty under which it can be guaranteed. These strategies respectively yield performance bands parameterized by presumed uncertainty level and required performance. Maximin performance bands for uncertainties in response-to-stimulus interval match data for the lower-scoring 70% of subjects well, and are more likely to explain it than robust-satisficing or alternative optimal performance curves. For uncertainties in signal-to-noise ratio, neither maximin nor robust-satisficing performance curves adequately describe the data. We discuss implications for decisions under uncertainties, and suggest further behavioral assays.

1 Introduction

There is a substantial literature on random walk and drift-diffusion (DD) processes as models for behavioral measures of reaction time and error rates on two-alternative forced-choice (2AFC) tasks, e.g. [1, 2, 3, 4, 5]. Additionally, *in vivo* recordings in monkeys trained to respond to motion stimuli show that neural spike rates in certain oculo-motor regions evolve like sample paths of a DD process, rising to a threshold prior to response initiation [6, 7, 8, 9, 10, 11]. Such experiments support a picture in which the state of the DD process, interpreted as a difference between accumulating, noisy evidence streams, is integrated until it reaches a confidence threshold.

The simple DD process with constant signal-to-noise ratio (SNR) is a continuum limit of the sequential probability ratio test (SPRT), which is optimal for statistically stationary tasks in the sense that, on average, it renders a decision of specified accuracy for the smallest number of observations [12, 13]. It therefore offers a normative theory of decision making against which human and animal behaviors can be assessed [14]. Specifically, both speed and accuracy are often at a premium, and a key issue is how these requirements are balanced in a *speed-accuracy tradeoff* (SAT). Wald and Wolfowitz [13] used a weighted sum of decision time and error rate as an objective function in their analysis of SPRT, and Edwards [15] generalized this to model how human subjects choose decision thresholds. This was further extended and tested in [16, 17], and a related theory, involving a competition between reward and accuracy (COBRA), has also been proposed [18, 19]. A review of these theories appears in [20], which also describes how the constant SNR DD process emerges from various evidence accumulator models, as parameters approach particular limits.

Gold and Shadlen [14] then proposed that the SPRT-DD model be used to derive an explicit SAT that optimizes reward rate alone, and Bogacz et al. [21, 20] performed this computation and carried out 2AFC experiments on human subjects to test its predictions. As discussed below, and at greater length in [22], the experiments reveal that, while a significant subset of 80 subjects performed near-optimally, others were substantially sub-optimal, favoring accuracy over speed.

Several explanations have been advanced to account for such behavior. Physiological constraints might prevent neural circuits from integrating evidence in the optimal manner prescribed by the SPRT (cf. [20] and see discussions below). Attention may wax and wane, leading to variable SNR (in fact, experimental protocols often mix stimuli of different discriminability, violating the stationarity assumed in the SPRT and in the blocked experiments of [22]). Intentional emphases on accuracy or speed, as in [18, 19], can modulate the SAT at the level of individual subjects. Finally, in addition to assuming statistical stationarity, the SPRT-based optimality theory of [14, 21, 20] implicitly assumes that key parameters such as experimenter-imposed response-to-stimulus delays and SNR are precisely known to the subject.

In this paper we shall briefly discuss all these factors before focusing on the last issue. Noting that optimal performance with normally distributed SNRs (e.g., due to variable drift rates) favors speed over accuracy [20] and thus cannot explain the observed SATs, we instead apply information gap theory [23] to predict SATs in the case that only uncertain estimates are available for key parameters. We find that

uncertainties in intertrial delays can account for sub-optimality in the data of [22], while uncertainties in the SNR cannot, and we show that this theory, which also involves only one free parameter, provides a better fit than a version of the COBRA theory that includes a weight for accuracy.

The paper is structured as follows. In §2 we review the DD model and derivation of the optimal SAT, describing how it can be expressed as an optimal performance curve (OPC). Information gap (Info-gap) extensions to the DD model, which account for parameter uncertainties, are developed in §3, and used to derive maximin and robust-satisficing¹ performance curves. The former selects thresholds that optimize performance for a given level of uncertainty; the latter predicts thresholds that maximize robustness for a pre-specified adequate performance level. In §4 we compare the predictions of the robust approaches with experimental data and with alternative optimization strategies. In §5 we compare the effects of parameter misestimates on the optimal and robust strategies, further assessing their predictions against the data of [22], and §6 contains a discussion and proposals for future experiments that may reveal which strategies are actually employed. Longer proofs and other mathematical and statistical details are relegated to Appendices. For general background on information gap decision theory, see [23].

2 A drift-diffusion model and optimal performance curves

The simplest version of a DD process is described by the following stochastic differential equation:

$$dx = A dt + \sigma dW ; \quad x(0) = 0, \quad (1)$$

where A denotes the drift rate and σdW increments drawn from a Wiener process with standard deviation σ . In the 2AFC context, the state variable $x(t)$ represents a running estimate of the logarithmic likelihood ratio [14], the two stimuli produce drift rates $\pm A$ respectively, and on each trial a choice is made when $x(t)$ first crosses either of the predetermined thresholds $\pm x_{th}$. It is implicitly assumed here that the stimuli are presented with equal probabilities, and, as in the optimal SPRT, $x(0)$ is initialized midway between the thresholds at the start of each trial and the SNR $|A/\sigma|$ remains constant for each block of trials.

Average performance on a block of trials with fixed SNR is characterized by the probability of making an error, $p(\text{err})$, and the *mean decision time*, $\langle DT \rangle$ (the first passage time for (1)), which can be computed explicitly as

$$p(\text{err}) = \frac{1}{1 + \exp(2\eta\theta)} \quad \text{and} \quad \langle DT \rangle = \theta \left[\frac{\exp(2\eta\theta) - 1}{\exp(2\eta\theta) + 1} \right], \quad (2)$$

[26, 27], cf. [20, Appendix], where the parameters $\eta = (A/\sigma)^2$ and $\theta = |x_{th}/A|$ are the square of the SNR (having the units of inverse time), and the threshold-to-drift ratio:

¹To satisfice is to accept a choice or judgment as one that is good enough. According to H. Simon [24], who coined the term, the tendency to satisfice appears in many cognitive tasks such as game playing, problem solving, and financial decisions, in which people typically do not or cannot seek optimal solutions [25]. We use the term here to distinguish the technical definition, given in §2, from the common notion of satisfying [23].

i.e., the passage time for the noise-free process $x(t) = At$. (The present notation differs from [20] to avoid confusion with that used below to describe uncertainty.) In (2) $\langle DT \rangle$ is that part of the *mean reaction time* $\langle RT \rangle$ ascribed to mental processing *per se*, excluding the non-decision latency T_0 required for sensory transduction and motor response, i.e., $\langle DT \rangle = \langle RT \rangle - T_0$.

The formulae of Eq. (2) may be inverted to solve for the DD model parameters, η and θ , in terms of the two performance parameters, $p(\text{err})$ and $\langle DT \rangle$:

$$\theta = \frac{\langle DT \rangle}{1 - 2p(\text{err})} \quad \text{and} \quad \eta = \frac{1 - 2p(\text{err})}{2\langle DT \rangle} \log \left(\frac{1 - p(\text{err})}{p(\text{err})} \right). \quad (3)$$

Given a suitable objective function for performance, these explicit formulae allow us to predict a parameter-free optimal performance curve that relates $p(\text{err})$ and $\langle DT \rangle$. Before describing this, we address some shortcomings of the simple DD model of Eq. (1).

2.1 Extensions to the DD model

As in [20] we shall refer to the constant SNR process of Eq. (1) as a *pure* DD process, to distinguish it from the *extended* processes widely studied by Ratcliff and others [3, 4, 5]. In view of the SPRT results noted above [12, 13], these generalized DD processes are necessarily non-optimal [20]. They were introduced to better account for human and animal behaviors.

First, it is observed that mean reaction times for error and correct trials typically differ, slow errors being frequently observed, whereas the pure DD process predicts identical forms for the probability distributions of DTs for both trial types. Slow errors can be produced by selecting the drift rate A for each trial from a Gaussian distribution, and fast errors result by selecting initial conditions $x(0)$ on each trial from a uniform distribution [4]. Including a possible overall bias, so that $\langle x(0) \rangle \neq 0$, these increase the number of DD parameters from two to five. To further allow for variability, the nondecision time T_0 can also be taken from a distribution [28]. Alternatively, time-dependent drift rates $A(t)$, which can represent varying attention, can produce slow or fast errors [29, 30], cf. [31, 32], in principle expanding the model to an infinite-dimensional parameter space.

Biophysical factors can be adduced to additionally complicate the model. Neurophysiology suggests that multiple populations of neurons, preferentially sensitive to different motion directions, are instrumental in decision making [6, 7, 8, 9, 10], and a model in wide use assumes a pair of leaky, competing accumulators [33] whose states represent averaged firing rates of two such neuronal pools. Since these are bounded above and below, this model involves nonlinear or piecewise-linear activation functions, and even if, as argued in [34], neural circuits self-tune to remain near their regions of maximal slope, the strengths of leak and inhibition are not necessarily equal. Only in this case can the linearized system be reduced to a pure DD process describing the difference between firing rates of the two pools [20], and recent simulations and studies of multi-unit spiking neuron models indicate that such a balance may be unlikely [35, 36]. Inclusion of a linear term λx in the drift

rate of (1) to represent the difference between inhibition and leak produces a linear Ornstein-Uhlenbeck (OU) process [26].

Unfortunately, for none of these generalized and non-optimal systems can first passage times and error rates be explicitly determined, although approximations for linear DD and OU processes are developed in [20]. Moreover, they all introduce additional parameters that must be fitted to individuals or groups of subjects. Variable drift-rates (either within or between trials) seem most relevant to account for differences in the reaction time distributions of error and correct decisions, but as noted in §2.2, the resulting optimal performance curves [20] fail to match the observed preference for accuracy over speed, which is the focus of this paper.

Here we wish to present a normative theory of *optimal* decision making, and to continue an investigation, begun in [20, 37] and [22] of how, and why, subjects diverge from its predictions. We do not intend to minimize the significant behavioral and physiological objections, described above, to the pure DD process, but rather to introduce a further factor, the degree of uncertainty in parameter estimates, and to contrast it with the emphasis placed on accuracy in theories such as COBRA [18, 19]. Furthermore, we show that uncertainties in SNR within the current framework also allow variable drift rates, and thus can account for differences in DT distributions for error and correct decisions. In summary, we contrast a parameter-free optimal SAT based on the pure DD process with several one-parameter generalizations that include weights for accuracy, levels of uncertainty and required performance. The predictions of these theories are then compared with data collected in experiments that, we believe, provide the opportunity and motivation for subjects to optimize a specific measure of performance.

2.2 Optimal performance curves

To determine optimal strategies for a given experimental paradigm, we must specify an objective or utility function. Here we consider the paradigm of [14], in which a subject is presented with a stimulus and is free to indicate a choice (by motor or oculomotor response) at any time. Correct responses are rewarded and after each response there is a fixed delay or response-to-stimulus interval (RSI), denoted D_{RSI} , before the next trial begins. This may be increased by a penalty delay $D_{pen} \geq 0$ in the event of an error. Trials are run sequentially in blocks in which delays and stimulus discriminabilities are constant, and whose durations are fixed, rather than the number of trials within them being predetermined, presenting subjects with a clear challenge of balancing accuracy with speed (the number of trials completed in the block).

Gold and Shadlen [14] propose that subjects seek to maximize their *average reward rates* (RR), defined as the fraction of correct responses divided by the average time between responses:

$$RR(p(\text{err}), \langle DT \rangle : D_{RSI}, T_0, D_p) = \frac{1 - p(\text{err})}{\langle DT \rangle + T_0 + D_{RSI} + p(\text{err})D_{pen}}. \quad (4)$$

Appendix A.1 shows that the mean number of correct decisions per unit time is given asymptotically by Eq. (4). For simplicity we have assumed that the non-decision

time T_0 and the RSI remain fixed during each block of trials over which the reward rate is to be maximized, but if these are variable one would simply replace them with their means $\langle T_0 \rangle$ and $\langle D_{RSI} \rangle$. Since only the combination $T_0 + D_{RSI}$ appears in (4), we define the non-decision delay $D = T_0 + D_{RSI}$, and for future use, the total delay $D_{tot} = T_0 + D_{RSI} + D_{pen}$.

Substituting Eqs. (2) into Eq. (4) yields the reward rate that can be achieved as a function of the DD parameters θ and η and the delays D and D_{tot} :

$$RR(\theta, \eta : D, D_{tot}) = [\theta + D + (D_{tot} - \theta) \exp(-2\eta\theta)]^{-1}. \quad (5)$$

The RSI D and penalty delay D_{pen} are under the experimenter's control, as is SNR $\sqrt{\eta}$ (at least partially), and if we assume that they and T_0 remain fixed over a block of trials, then the threshold-to-drift ratio θ is the only adjustable parameter and local maxima of RR occur at zeros of $\partial RR / \partial \theta$:

$$\exp(2\eta\theta) - 1 = 2\eta(D_{tot} - \theta). \quad (6)$$

The left hand side of (6) is monotonically increasing with θ and its right hand side is monotonically decreasing, so it has a unique solution $\theta = \theta_{op}$ that depends only upon η and D_{tot} , and that determines the globally optimal normalized threshold. The corresponding optimal reward rate RR_{op} , which also depends on D , is:

$$RR_{op}(\eta : D, D_{tot}) = [\theta_{op} + D + (D_{tot} - \theta_{op}) \exp(-2\eta\theta_{op})]^{-1}. \quad (7)$$

Eq. (6) establishes a relationship among θ , η and D_{tot} that must hold if RR is maximized, and we may use Eqs. (3) to replace the first two DD model parameters by the performance parameters $\langle DT \rangle$ and $p(\text{err})$, to get:

$$\frac{\langle DT \rangle}{D_{tot}} = \left[\frac{1}{p(\text{err}) \log \left(\frac{1-p(\text{err})}{p(\text{err})} \right)} + \frac{1}{1 - 2p(\text{err})} \right]^{-1}. \quad (8)$$

This *optimal performance curve* (OPC), first derived in [20], is a unique, parameter-free relationship that describes the tradeoff between speed ($\langle DT \rangle / D_{tot}$) and accuracy ($1 - p(\text{err})$) that must hold in order to maximize RR . The curve limits at zero at both ends of the admissible range of error rates and has a maximum at $p(\text{err}) \approx 0.1741$ (see below). It is important to understand that *all points on the OPC (all $(p(\text{err}), \langle DT \rangle / D_{tot})$ pairs) correspond to optimal performances*, under different conditions.

To better appreciate this we note that a subject's behavioral location in the $(p(\text{err}), \langle DT \rangle / D_{tot})$ -plane on a given block of trials can be manipulated either by varying stimulus discriminability (which partially determines SNR η), or by changing the total delay D_{tot} , for example via the RSI D . The form of the OPC may be intuitively understood by observing that very noisy stimuli ($\eta \approx 0$) contain little information, so that given the a priori knowledge that they are equally likely, it is optimal to choose at random without examining them; thus $p(\text{err}) = 0.5$ and $\langle DT \rangle = 0$. At the opposite limit ($\eta \rightarrow \infty$), noise-free stimuli are so easy to discriminate




Figure 1: Thick curve shows the optimal performance curve of Eq. (6), and histograms show data collected from 80 human subjects, sorted according to total rewards accrued over multiple blocks of trials with two difficulty levels and with $D_{RSI} = 0.5, 1.0, 2.0$ s and $D_p = 0, 1.5$ s. White bars: all subjects; lightest bars: lowest 10% excluded; medium bars: lowest 50% excluded; darkest bars: lowest 70% excluded. Vertical line segments indicate standard errors. From [21, Fig. 1].

that $\langle DT \rangle$ and $p(\text{err})$ both approach zero. In all other cases, it is advantageous to accumulate the noisy evidence for long enough to make an informed choice, and the height of the curve predicts the optimal accumulation time. Since the OPC is independent of the parameters defining the DD process, and D_{tot} enters only as the denominator of the normalized decision time, data from different subjects and blocks of trials, with differing SNR and D_{tot} , can be combined in a single histogram for comparison with the OPC.

To test whether humans can optimize their performance in this sense, two experiments were performed under the above paradigm, as detailed in [20, 22]. In the first, 20 subjects viewed motion stimuli [38] and received 1 cent for each correct discrimination. The experiment was divided into 7-minute blocks with different RSIs and penalty delays: three with $D_p = 0$ and $D_{RSI} = 0.5\text{s}, 1.0\text{s},$ and 2.0s , respectively, and one with $D_p = 1.5\text{s}$ and $D_{RSI} = 0.5\text{s}$. One block of trials was run for each condition except for $D_p = 0, D_{RSI} = 2.0\text{s}$, for which two blocks were required to gather sufficient data. In the second experiment, 60 subjects discriminated if the majority of 100 locations on a visual display were filled with stars or empty [4]. The design was the same as that of the first experiment except that blocks lasted for 4 minutes, and the set of 5 blocks was repeated under two difficulty conditions. In both cases subjects were instructed to maximize their total earnings, and unrewarded practice blocks were administered before the tests began.

For each subject and condition $p(\text{err})$ s were estimated by the error rates and $\langle DT \rangle$ s were estimated by fitting the DD model to experimental reaction time distributions, as described in [20]. Data from both experiments were pooled and sorted by error rates into 5% bins, and the resulting mean of the estimated $\langle DT \rangle$ normalized by D_{tot} are shown by white bars in Fig. 1. The shaded bars show results of the same analysis restricted to subgroups of subjects ranked according to their total rewards accrued over all blocks of trials and experimental conditions. This indicates that the OPC matches performances of the top 30% of subjects, who appear to choose their SATs to optimize rewards under each condition (the darkest bars lie close to the OPC). However, it clearly fails to capture the behaviors of subgroups with lower total scores, who typically exhibit longer $\langle DT \rangle$'s, implicitly valuing accuracy over speed.

We believe that the conditions under which this data was collected, with stimulus discriminability and RSIs fixed during each block of trials, along with training sessions and emphasis on maximizing rewards, strongly encouraged subjects to improve their ability to extract signal from noise, thereby driving their decision mechanisms closer to that described by the optimal DD process. Indeed, while fits to an extended DD process with variability in A and $x(0)$ improved on those to the pure DD process, detailed comparison, summarized in Appendix D, reveals that the two models capture a similar proportion of the variance in subjects' thresholds. Furthermore, allowing variability in A and $x(0)$ results in OPCs that differ qualitatively from the experimental data [20, Fig. 14]. Specifically, trial-to-trial variability in A decreases the mean normalized decision time for a given error rate, and variability in $x(0)$ results in higher than observed mean normalized decision time at zero error rate. We therefore believe that it is reasonable to use pure DD fits in comparing the data from all 80 subjects with the OPC and modified OPCs derived from the pure DD model.

2.3 Performance curves weighted for accuracy

The tendency to value accuracy over speed has been noted earlier [39, 18, 19], and it has suggested alternative objective functions [21, 20]. Here we describe two examples. The first is a modified *reward-accuracy rate* (RA) that additionally penalizes errors, as suggested by the COBRA theory [18, 19]:

$$RA = RR - \frac{q p(\text{err})}{D_{tot}}. \quad (9)$$

Here q specifies the relative weight placed on accuracy, and the characteristic time D_{tot} is included so that the units of both terms are $(\text{time})^{-1}$. Maximizing the expression in (9), we obtain a family of OPCs parameterized by q :

$$\frac{\langle DT \rangle}{D_{tot}} = \frac{\mathcal{E} - 2q - \sqrt{\mathcal{E}^2 - 4q(\mathcal{E} + 1)}}{2q}, \quad (10)$$

$$\text{where } \mathcal{E} = \left[\frac{1}{p(\text{err}) \log \left(\frac{1-p(\text{err})}{p(\text{err})} \right)} + \frac{1}{1 - 2p(\text{err})} \right] \quad (11)$$

is the reciprocal of the OPC formula (8). In Appendix A.2 we prove that \mathcal{E} has a unique minimum (and hence the OPC has a unique maximum) at $p(\text{err}) \approx 0.1741$. It follows that the nondimensional decision times given by Eq. (10) are non-negative real numbers provided that the weight satisfies the following inequality:

$$q \leq \min \left\{ \frac{\mathcal{E}^2}{4(\mathcal{E} + 1)} \right\} = \frac{\mathcal{E}_{min}^2}{4(\mathcal{E}_{min} + 1)}. \quad (12)$$

Numerically, we find $\mathcal{E}_{min} \approx 5.224$, implying that $q \leq 1.096$. Moreover, differentiating the expression in (10) with respect to $p(\text{err})$, we find that its critical points coincide with the minimum of \mathcal{E} for all q satisfying (12), implying that the maxima of the OPC for RA all occur at $p(\text{err}) \approx 0.1741$, as for RR .

Secondly, for monetary rewards one can suppose that errors are penalized by subtraction from previous winnings, leading to the modified reward rate:

$$RR_m = \frac{(1 - p(\text{err})) - qp(\text{err})}{\langle DT \rangle + D_{tot}}, \quad (13)$$

which leads to the following OPC family:

$$\frac{\langle DT \rangle}{D_{tot}} = (1 + q) \left[\frac{\frac{1}{p(\text{err})} - \frac{q}{1-p(\text{err})}}{\log \left(\frac{1-p(\text{err})}{p(\text{err})} \right)} + \frac{1 - q}{1 - 2p(\text{err})} \right]^{-1}. \quad (14)$$

Unlike the OPCs for RA , the maxima of the family (14) move rightward with increasing q : see Fig. 6 below. The nondimensional decision times given by Eq. (14) are non-negative provided that the weight satisfies $q \leq 1$.

The utility functions defining RA and RR_m each involve a single weight parameter q , and Eqs. (10) and (14) both reduce to (8) for $q = 0$. Since errors are explicitly penalized in the expressions (9) and (13), no additional delay D_p is included in D_{tot} here. In §4 we shall assess these modified OPCs against the data, along with predictions from information gap theory (Fig. 6).

In the next section we develop decision strategies for the 2AFC task that are robust against uncertainties in either delays (§3.1) or SNR (§3.2). We show that they differ from the optimal strategies described above, and in §4 we argue that the data is more consistent with the use of robust decision strategies against uncertainties in delays than for any of the optimal procedures, especially for the lower 70% of subjects based on total scores. In common with optimal ones, robust strategies require threshold setting, or, equivalently, the formulation of a stopping criterion for the decision procedure. We do not consider these issues here, but see [39, 40] for models of criterion learning, and see [37] for a neural network implementation of threshold setting in the DD model context.

3 Robust decisions under uncertainties

In optimal decision theory subjects are supposed to maximize a utility function (RR in the 2AFC task) for the actual task parameters. This assumption is questionable,

given typical constraints on time, knowledge and computational capacity [41, 24, 42, 43]. In particular, it has been argued that people do not have unlimited knowledge to base decisions upon, and that optimal strategies may not result in optimal outcomes due to simplifying assumptions. As noted in §2, a full model of the 2AFC process may be complicated, and neither its exact structure nor its parameters well known.

Simon [41, 24] proposed that people may instead satisfice an aspiration level, rather than optimize performance. A simple satisficing strategy is to choose the first alternative that meets or exceeds the aspiration level. When all alternatives are known, as in the case of 2AFC, the main search is for the proper cue to evaluate them. The Take-The-Best heuristic proposed in [44, 43], for example, suggests focusing on the best cue that differentiates between the alternatives. In the context of DDM, this may suggest using the drift rate as the cue; but the issue of determining the threshold is left open.

Considering the robustness of the satisficing strategy, Ben-Haim [23] suggested choosing the alternative that achieves the aspiration level (or required performance) under the largest level of uncertainty. This *robust-satisficing* strategy can explain decision making paradoxes such as that of Ellsberg [45] by positing relatively high levels of desired robustness for success, and it has been applied successfully to explain foraging behavior, where a critical level of food intake is necessary for survival [46], and equity premium, in which a critical level of return is required [23]. In contrast with this, the *maximin* decision strategy selects the alternative that optimizes the worst performance given a presumed level of uncertainty. We shall develop both maximin and robust-satisficing strategies for 2AFC tasks under uncertainties, and investigate how well they explain the observed SATs.

There are at least four sources of uncertainty in using the DD process of Eq. (1) to model 2AFC experiments: (i) uncertainties in delays and SNR; (ii) uncertainties in the level of the threshold enforced by the neural system; (iii) uncertainties in the distribution of noise, and (iv) uncertainties in the structure of the underlying model. Here we address only the first, assuming that choices are indeed described by the DD model specified by Eq. (1) with white noise, zero initial condition and appropriately chosen symmetric thresholds. We consider uncertainties in delays in §3.1 and uncertainties in SNR, due to either noise variance or drift rate, in §3.2. The latter analysis provides an alternative approach to drawing from a distribution of drift rates on each trial, as in the extended DD model. Here the drift rates on each trial may vary within a bounded range, but not necessarily according to any probabilistic distribution.

3.1 Uncertainties in delays

The uncertainty model: Uncertainties in delays may arise from (i) poor estimates of RSIs: the scalar property of interval timing implies that estimated delays are distributed around the actual value with a standard deviation proportional to it [47, 48], and (ii) uncertainties in the non-decision latency T_0 due to trial-to-trial variations in sensory processing and motor response. Following Eq. (5), we consider uncertainties in both $D = D_{RSI} + T_0$ and $D_{tot} = D + D_p$. The best estimates or

most likely values are termed the *nominal delays* and denoted by: $\tilde{D} = \tilde{D}_{RSI} + \tilde{T}_0$ and $\tilde{D}_{tot} = \tilde{D} + \tilde{D}_p$.

Since their actual values are not necessarily random variables, we characterize the uncertainties by a non-probabilistic information gap (*Info-gap*) model, which specifies a nested, unbounded set of intervals for the actual values, parameterized by the level of uncertainty. Nesting implies that intervals associated with larger levels of uncertainty include intervals associated with smaller levels of uncertainty, and all the intervals include the nominal values, which are associated with zero uncertainty. Unboundedness implies that the intervals increase without bound as the level of uncertainty increases [23]. In agreement with the scalar property of interval timing [47, 48] we postulate that the size of the intervals is proportional to the nominal delays according to a proportional Info-gap model [23]:

$$U(\alpha, \tilde{D}, \tilde{D}_{tot}) = \left\{ D > 0, D_{tot} > 0 : \left| D - \tilde{D} \right| \leq \alpha \tilde{D}, \quad \left| D_{tot} - \tilde{D}_{tot} \right| \leq \alpha \nu \tilde{D}_{tot} \right\}. \quad (15)$$

These intervals are parameterized by the overall uncertainty $\alpha \geq 0$ and the relative weight ν between uncertainties in the different delays. Henceforth we set $\nu = 1$ for simplicity.

Maximin decision strategy: As the level of uncertainty increases, deviations between the actual and nominal delays may grow. Performance as measured by RR, which depends on the delays (Eq. (5)), decreases under unfavorable delays and increases under favorable ones, even when the threshold and SNR are fixed. The optimal strategy of §2.2 maximizes the RR that can be achieved with the nominal delays, but ignores potential degradation in performance due to unfavorable delays. The maximin strategy instead focuses on the minimal RR that can be achieved given a presumed level of uncertainty α_p , and then selects the threshold θ_{MM} that maximizes this level.

Definition 3.1. Maximin threshold under uncertainties in delays: *The maximin threshold θ_{MM} is the threshold that maximizes the worst RR under the presumed level of uncertainty α_p in the delays, given the SNR and nominal delays:*

$$\theta_{MM}(\alpha_p : \eta, \tilde{D}, \tilde{D}_{tot}) = \arg \max_{\theta} \left(\min_{D, D_{tot} \in U(\alpha_p, \tilde{D}, \tilde{D}_{tot})} RR \right). \quad (16)$$

Theorem 3.1. Maximin threshold under uncertainties in delays: *Given the presumed level of uncertainty α_p in the Info-gap model of the delays specified by Eq. (15), the maximin threshold θ_{MM} satisfies:*

$$\exp(2\eta\theta_{MM}) - 1 = 2\eta[\tilde{D}_{tot}(1 + \alpha_p) - \theta_{MM}]. \quad (17)$$

Proof: The inner minimization in Eq. (16) specifies the worst performance when the level of uncertainty in the delays is α_p . In Appendix B.1 we derive the worst performance θ_{MM} under uncertainties in the delays that are described by the Info-gap

model of Eq. (15). It is shown that the maximum *inverse reward rate* I_{MAX} (Eq. (44) with $\alpha = \alpha_p$) is given by:

$$I_{MAX}(\alpha_p : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) = \theta + \tilde{D}(1 + \alpha_p) + [\tilde{D}_{tot}(1 + \alpha_p) - \theta] \exp(-2\eta\theta) \quad (18)$$

Thus, θ_{MM} is the threshold that minimizes Eq. (18), and hence nullifies its derivative with respect to θ . Eq. (17) is obtained by differentiating Eq. (18) with respect to θ , and setting the derivative to zero. It is straightforward to verify that this is a minimum by evaluating the second derivative at that point. \square

When the presumed robustness is zero, Eq. (17) reduces to the optimum threshold expression (6).

Substituting Eqs. (3) in Eq. (17), we can derive a *maximin performance curve* (MMPC), which describes the tradeoff between speed ($\langle DT \rangle / \tilde{D}_{tot}$) and accuracy ($1-p(\text{err})$) that must hold in order to maximize the worst *RR* under a fixed level of presumed uncertainty α_p in the delays:

$$\frac{\langle DT \rangle}{\tilde{D}_{tot}} = \gamma \left[\frac{1}{p(\text{err}) \log \left(\frac{1-p(\text{err})}{p(\text{err})} \right)} + \frac{1}{1-2p(\text{err})} \right]^{-1}, \quad (19)$$

where $\gamma = 1 + \alpha_p$.

For comparison with experimental data, the nominal delay \tilde{D}_{tot} must be related to the actual total delay D_{tot} . Assuming equality, i.e., $D_{tot} = \tilde{D}_{tot}$, defines a one-parameter family of nominal MMPCs:

$$\frac{\langle DT \rangle}{D_{tot}} = \gamma \left[\frac{1}{p(\text{err}) \log \left(\frac{1-p(\text{err})}{p(\text{err})} \right)} + \frac{1}{1-2p(\text{err})} \right]^{-1}. \quad (20)$$

When actual delays differ from nominal values, different MMPCs are obtained. If actual delays vary within a given range, the performance lies within a corresponding *performance band*, which is bounded by the MMPCs for the shortest and longest delays. Considering, for example, the case when delays vary within the range consistent with the presumed uncertainty, i.e., $D_{tot} \in U(\alpha_p, \tilde{D}, \tilde{D}_{tot})$; the MMPC associated with the longest delay $D_{tot} = \tilde{D}_{tot}(1 + \alpha_p)$ coincides with the OPC, while the extreme MMPC associated with the shortest delay $D_{tot} = \tilde{D}_{tot}(1 - \alpha_p)$ has the same form as Eq. (20) with $\gamma = (1 + \alpha_p)/(1 - \alpha_p)$.

Fig. 2 shows MMPCs and performance bands for two presumed levels of uncertainty. As indicated by Eq. (20), the MMPCs for different α_p are simply scaled versions of the OPC (Eq. (8)), which is the special case of an MMPC with zero presumed uncertainty ($\gamma = 1$). Thus, all the MMPCs peak at the same $p(\text{err}) \approx 0.1741$. For a fixed α_p , the maximin strategy enforces performance to remain along the corresponding MMPC defined by Eq. (20). Each point along a MMPC corresponds to a different θ_{MM} , which satisfies the maximin strategy for specific SNR and nominal delays, as defined by Eq. (16) and specified by Eq. (17).




Figure 2: Maximin performance curves (MMPCs) and performance bands for two presumed uncertainty levels $\alpha_p = 0.2, 0.4$ in delays. Nominal MMPCs for $D_{tot} = \tilde{D}_{tot}$ (solid), and extreme MMPCs for $D_{tot} = \tilde{D}_{tot}(1 - \alpha_p)$ (dashed) and for $D_{tot} = \tilde{D}_{tot}(1 + \alpha_p)$ (dot-dashed) are shown; the latter coincide with the OPC for all α_p (cf. Eq. (19) with $D_{tot} = \tilde{D}_{tot}(1 + \alpha_p)$ and Eq. (8)). Performance bands for delays lying within the range consistent with the presumed uncertainty are bounded by the extreme MMPC and the OPC.

Fig 2 demonstrates that the MMPCs and bands capture the qualitative form of the experimental data fairly well, especially (i) the higher $\langle DT \rangle$'s, and (ii) the wide spread in $\langle DT \rangle$'s for error rates in the range 15-35%. Quantitative comparisons are given in §4.

Robust-satisficing decision strategy: The maximin strategy assumes a fixed level of uncertainty, which may not be known. Furthermore, it appeals to worst case performance, which may be very unlikely, and hence is a prudent or pessimistic strategy. Instead robust-satisficing focuses on the required performance without making any assumptions about the level of uncertainty.

Definition 3.2. Robustness to uncertainties in delays: *Consider a 2AFC strategy with threshold θ and a required performance R_r . The robustness $\hat{\alpha}$ to uncertainties in the delays is the greatest horizon of uncertainty α in the Info-gap model of the delays for which a required performance level R_r can be guaranteed. Robustness is zero if the required performance cannot be assured even without uncertainty in delays:*

$$\hat{\alpha}(R_r : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) = \max \left\{ \alpha \left| \left(\min_{D, D_{tot} \in U(\alpha, \tilde{D}, \tilde{D}_{tot})} RR \right) \geq R_r \vee \alpha = 0 \right. \right\}. \quad (21)$$

The functional dependence of the robustness on the required performance is termed the robustness curve.

Robustness curves associated with the Info-gap model of Eq. (15), are derived in Appendix B.2, Eq. (46), and depicted there in Fig 9. Note that the nominal performance, which can be achieved with the nominal delays, $RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot})$, has zero robustness. In particular, the nominal optimal performance of Eq. (7), $RR_{op}(\eta : \tilde{D}, \tilde{D}_{tot})$, achieved with the optimal threshold θ_{op} under the nominal delays, has zero robustness. Only lower reward rates $R_r < RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot}) \leq RR_{op}(\eta : \tilde{D}, \tilde{D}_{tot})$, can be attained robustly. Robustness to uncertainties requires relinquishing performance.

Definition 3.3. Robust-satisficing threshold under uncertainties in delays: *The robust-satisficing threshold θ_{RS} corresponding to the required sub-optimal performance R_r is the threshold that maximizes robustness to uncertainties in delays:*

$$\theta_{RS}(R_r : \eta, \tilde{D}, \tilde{D}_{tot}) = \arg \max_{\theta} \hat{\alpha}(R_r : \theta, \eta, \tilde{D}, \tilde{D}_{tot}). \quad (22)$$

The robust-satisficing strategy assures the required sub-optimal performance R_r under the largest level of uncertainty. While the optimal nominal performance has zero robustness, sub-optimal performance levels may be attained with positive robustness to uncertainties in the delays, as specified in Eq. (21). The robust-satisficing strategy seeks the threshold θ_{RS} that provides maximum robustness; at the expense of relinquishing the optimal level that could be achieved when there are no uncertainties.

Theorem 3.2. Robust-satisficing threshold under uncertainties in delays: *Consider a 2AFC strategy with uncertainties in the delays specified by the Info-gap model of Eq. (15). Given a sub-optimal required performance $R_r < RR_{op}(\eta : \tilde{D}, \tilde{D}_{tot})$, the robust satisficing threshold θ_{RS} , which can achieve the required performance with maximum robustness to uncertainties in delays, satisfies:*

$$2\eta\tilde{D}_{tot}I_r = [2\eta\theta_{RS} + \exp(2\eta\theta_{RS}) - 1] \tilde{D} + [2\eta\theta_{RS} - \exp(-2\eta\theta_{RS}) + 1] \tilde{D}_{tot}. \quad (23)$$

Proof: See Appendix B.3. □

Substituting Eq. (3) in Eq. (23), we derive the *robust-satisficing performance curve* (RSPC) that describes the tradeoff between speed ($\langle DT \rangle / \tilde{D}_{tot}$) and accuracy ($1-p(\text{err})$) that must hold in order to robust-satisfice the required performance R_r :

$$\frac{\langle DT \rangle}{\tilde{D}_{tot}} = \frac{R_r^{-1}}{\tilde{D}_{tot}} \left[\frac{\tilde{D}/\tilde{D}_{tot} + p(\text{err})(1 - \tilde{D}/\tilde{D}_{tot})}{p(\text{err})(1 - p(\text{err})) \log\left(\frac{1-p(\text{err})}{p(\text{err})}\right)} + \frac{\tilde{D}/\tilde{D}_{tot} + 1}{1 - 2p(\text{err})} \right]^{-1}. \quad (24)$$

Note that, even when the actual delays equal the nominal values ($D = \tilde{D}$ and $D_{tot} = \tilde{D}_{tot}$), or when there is no penalty delay ($\tilde{D}_{tot} = \tilde{D}$), this expression does not reduce to a scaled OPC (8). The RSPCs have a different shape.

RSPCs depend on the ratio $\tilde{D}/\tilde{D}_{tot}$, so for comparison with experimental data, this ratio must be determined or related to the actual ratio D/D_{tot} . When $D_p = 0$,

the ratio is one, so the RSPCs can be computed from Eq. (24), by substituting $\tilde{D}/\tilde{D}_{tot} = 1$:

$$\frac{\langle DT \rangle}{D_{tot}} = \frac{R_r^{-1}}{D_{tot}} \left[\frac{1}{p(\text{err})(1-p(\text{err})) \log\left(\frac{1-p(\text{err})}{p(\text{err})}\right)} + \frac{2}{1-2p(\text{err})} \right]^{-1}. \quad (25)$$

Fig. 3 depicts such RSPCs for three different levels of normalized required performance $q \equiv (R_r^{-1} - D)/D_{tot}$. The RSPCs all peak at $p(\text{err}) \approx 0.1352$: a leftward shift compared with the OPCs and MMPCs that peak at $p(\text{err}) \approx 0.1741$. Since the RSPCs are derived for the case $D_{pen} = 0$, they must be evaluated against experimental results with zero penalty delay. This is deferred to §4, where it is shown that the leftward shift slightly degrades the fit with the data (see Fig. 6 and Table 3).




Figure 3: Robust satisfying performance curves (RSPCs, solid) for uncertainties in delays given three normalized required performance levels $q = (R_r^{-1} - D)/D_{tot}$. Here the nominal delay ratio $\tilde{D}/\tilde{D}_{tot} = 1$ is assumed, corresponding to the case of no delay penalty. The OPC, depicted by the dot-dashed line, constrains the RSPCs below. The dashed line below the OPC for $q = 0.5$ extends the RSPC into the opportunity facilitating performance curve (OFPC).

Each RSPC (Eq. (24)) describes the speed-accuracy tradeoff that results from adopting the robust satisfying strategy for a fixed R_r . Each point along an RSPC corresponds to a different θ_{RS} , which satisfies the robust satisfying strategy for specific SNR and nominal delays, as defined by Eq. (22) and specified by Eq. (23). However, as SNRs degrade and delays increase, the optimal performance RR_{op} (Eq. (7)) becomes smaller and eventually falls below the required performance R_r . Since the robust satisfying strategy is relevant only for sub-optimal required performance $R_r < RR_{op}(\eta : \tilde{D}, \tilde{D}_{tot})$, the RSPCs are constrained to lie above the OPC (see also Appendix B.3).

While better than optimal performance cannot be guaranteed, it may nevertheless occur under favorable conditions. As indicated above, the robust-satisficing

strategy is irrelevant for such *windfall performance levels* (denoted by R_w). Instead, Info-gap theory [23] suggests an alternative strategy that facilitates the pursuit of windfall performance, as detailed in Appendix B.4. The resulting *opportunity facilitating performance curves* (OFPCs) extend the RSPCs given by Eq. (24) for $R_w > RR_{op}(\eta : \bar{D}, \bar{D}_{tot})$ as depicted by the dashed curve for $q = 0.5$ in Fig. 3.

3.2 Uncertainties in signal-to-noise ratio

We now consider robust decisions to uncertainties in SNR ($\sqrt{\eta}$), assuming that delays are known. Such uncertainties may arise from poor estimates of a fixed SNR, or from trial-to-trial variations, and we further differentiate between uncertainties in drift rate and uncertainties in variance of the noise. The source of uncertainties is important since the threshold-to-drift ratio θ depends on drift rate but not on noise. Uncertainties due to noise variance are considered in §3.2.1 and those due to uncertainties in drift in §3.2.2. In both cases the SNR is assumed fixed, but as noted in §2.2, trial-to-trial variations in drift rate have been invoked to explain the observed differences between correct and error response times. In §3.2.3 we therefore extend our analysis of uncertainties in drift to include such variations.

As in §3.1, a proportional Info-gap model is formulated and parameterized by $\alpha \geq 0$:

$$U(\alpha, \tilde{\eta}) = \{\eta > 0 : |\eta - \tilde{\eta}| \leq \alpha \tilde{\eta}\}. \quad (26)$$

where $\sqrt{\tilde{\eta}}$ is the nominal SNR.

3.2.1 Uncertainties in noise variance

Here we assume that in each experiment the drift rate A is fixed and known. Hence, determining the threshold x_{th} is equivalent to determining the threshold-to-drift ratio θ .

Definition 3.4. Maximin threshold-to-drift ratio under uncertainties in the SNR: *The maximin threshold-to-drift ratio θ_{MM} is the one that maximizes the worst RR under the presumed level of uncertainty α_p in the SNR, given the delays and nominal SNR:*

$$\theta_{MM}(\alpha_p : \tilde{\eta}, D, D_{tot}) = \arg \max_{\theta} \left(\min_{\eta \in U(\alpha_p, \tilde{\eta})} RR \right). \quad (27)$$

Theorem 3.3. Maximin threshold-to-drift ratio under uncertainties in SNR: *Given the presumed level of uncertainty $\alpha_p \leq 1$ in the Info-gap model of the SNR given by Eq. (26), the maximin threshold-to-drift θ_{MM} satisfies:*

$$\exp(2\tilde{\eta}(1 - \alpha_p)\theta_{MM}) - 1 = 2\tilde{\eta}(1 - \alpha_p)(D_{tot} - \theta_{MM}), \quad (28)$$

Proof: The inner minimization in Eq. (27) specifies the worst performance when the level of uncertainty in the SNR is α_p . In Appendix B.5 we derive the worst performance under uncertainties in the SNR that are described by the Info-gap

model of Eq. (26). It is shown that the maximum inverse reward rate I_{MAX} (Eq. (56) with $\alpha = \alpha_p$) is given by:

$$I_{MAX}(\alpha_p \leq 1 : \theta, \tilde{\eta}, D, D_{tot}) = \theta + D + (D_{tot} - \theta) \exp(-2\tilde{\eta}(1 - \alpha_p)\theta). \quad (29)$$

Thus, θ_{MM} is the threshold that minimizes Eq. (29), and hence nullifies its derivative with respect to θ . Eq. (28) is obtained by differentiating Eq. (29) with respect to θ , and setting the derivative to zero. It is straightforward to verify that this is a minimum by evaluating the second derivative at that point. \square

When presumed robustness is zero, Eq. (28) for the maximin threshold-to-drift ratio reduces to Eq. (6) for the optimum threshold.

Assuming that the nominal SNR is the actual SNR ($\eta = \tilde{\eta}$), the nominal MMPC for a fixed level of presumed uncertainty $\alpha_p \leq 1$ in the SNR, can be derived by substituting Eq. (3) in Eq. (28):

$$\frac{\langle DT \rangle}{D_{tot}} = [1 - 2p(\text{err})] \left[\frac{\left(\frac{1-p(\text{err})}{p(\text{err})} \right)^\chi - 1}{\chi \log \left(\frac{1-p(\text{err})}{p(\text{err})} \right)} + 1 \right]^{-1}, \quad (30)$$

where $\chi \equiv 1 - \alpha_p$. Under uncertainties in the SNR, the MMPCs are not scaled version of the OPC (Eq. (8)), but reduce to it when the uncertainty vanishes ($\chi = 1$).

If the SNR varies within the range consistent with the presumed level of uncertainty, i.e., $\eta \in [\tilde{\eta}(1 - \alpha_p), \tilde{\eta}(1 + \alpha_p)]$ for $\alpha_p \leq 1$, the maximin performance bands are bounded below (for $\eta = \tilde{\eta}(1 - \alpha_p)$) by the OPC (Eq. (8)), and above (for $\eta = \tilde{\eta}(1 + \alpha_p)$) by Eq. (30) with $\chi = \frac{1 - \alpha_p}{1 + \alpha_p}$. Fig. 4 shows MMPCs and performance bands for two levels of presumed uncertainty. The curves exhibit a leftward shift in peak values that is not characteristic of the data of Fig. 1, and in contrast with MMPCs for uncertainties in delays of Fig. 2, lack the higher $\langle DT \rangle$'s at mid error rates. We shall nonetheless consider them further in §4, in order to quantitatively assess the relative likelihood of the two sources of uncertainty in accounting for suboptimal performance.

Robust satisficing and opportunity facilitating: Under uncertainties in SNR that arise from uncertainties in the noise variance, the robust satisficing and opportunity facilitating strategies result in linear performance curves, as detailed in Appendix B.6. These RSPCs start at finite $\langle DT \rangle$ at $p(\text{err}) = 0$ and decline to zero at $p(\text{err}) = 0.5$ as shown in Fig. 10, and they also fail to capture the higher $\langle DT \rangle$'s for mid $p(\text{err})$'s apparent in Fig. 1. Consequently, these strategies are not considered in the comparisons of §4.

3.2.2 Uncertainties in drift rate

We now consider uncertain drift rates. Uncertainties in A affect the threshold-to-drift ratio $\theta = x_{th}/A$, so this parameter is not suitable for determining the threshold under such uncertainties. We replace it by the threshold-to-noise ratio $\vartheta = x_{th}/\sigma$, and, since σ^2 is assumed fixed and known, determining the threshold x_{th} is now




Figure 4: MMPCs for two presumed levels $\alpha_p = 0.3, 0.6$ of uncertainties in SNR arising from uncertain noise variance. Nominal MMPC for $\eta = \tilde{\eta}$ (solid), and extreme MMPC for $\eta = \tilde{\eta}(1 + \alpha_p)$ (dashed) and for $\eta = \tilde{\eta}(1 - \alpha_p)$ (dot-dashed); the latter coincide with the OPC for all α_p . Performance bands for SNR within the range consistent with the presumed uncertainty are bounded by the extreme MMPC and the OPC.

equivalent to determining ϑ . Note that $\vartheta = \theta\sqrt{\eta}$, so the quantities describing performance: $p(\text{err})$ and $\langle DR \rangle$ in Eq. (2) and RR in Eq. (5), can be expressed in terms of ϑ and η in place of θ and η .

Definition 3.5. Maximin threshold-to-noise ratio under uncertainties in SNR: *The maximin threshold-to-noise ratio ϑ_{MM} is the one that maximizes the worst RR , under the presumed level of uncertainty α_p in the SNR, given the delays and nominal SNR:*

$$\vartheta_{MM}(\alpha_p : \tilde{\eta}, D, D_{tot}) = \arg \max_{\vartheta} \left(\min_{\eta \in U(\alpha_p, \tilde{\eta})} RR \right). \quad (31)$$

Theorem 3.4. Maximin threshold-to-noise ratio under uncertainties in SNR: *Given the presumed level of uncertainty $\alpha_p \leq 1$ in the Info-gap model of the SNR given by Eq. (26), the maximin threshold-to-noise ϑ_{MM} satisfies:*

$$\exp(2\sqrt{\tilde{\eta}(1 - \alpha_p)\vartheta_{MM}}) - 1 = 2\tilde{\eta}(1 - \alpha_p)(D_{tot} - \frac{\vartheta_{MM}}{\sqrt{\tilde{\eta}(1 - \alpha_p)}}), \quad (32)$$

Proof: See Appendix B.7. □

Given $\tilde{\eta}$, we may substitute $\theta_{MM} = \vartheta_{MM}/\sqrt{\tilde{\eta}}$. Assuming that the nominal SNR is the actual SNR ($\eta = \tilde{\eta}$), the nominal MMPC for a fixed level of presumed

uncertainty $\alpha_p \leq 1$ in the SNR, can be derived by substituting Eq. (3) in Eq. (32):

$$\frac{\langle DT \rangle}{D_{tot}} = \sqrt{\chi} (1 - 2p(\text{err})) \left[\frac{\left(\frac{1-p(\text{err})}{p(\text{err})} \right)^{\sqrt{\chi}} - 1}{\sqrt{\chi} \log \left(\frac{1-p(\text{err})}{p(\text{err})} \right)} + 1 \right]^{-1}, \quad (33)$$

where, as in Eq. (30), $\chi \equiv 1 - \alpha_p$.

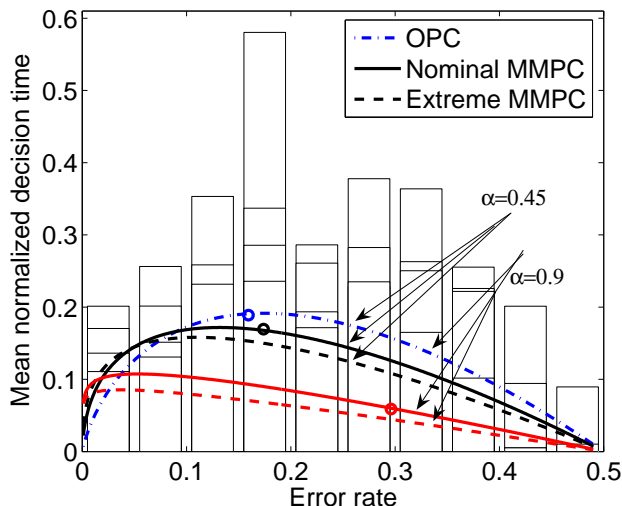


Figure 5: MMPCs for two presumed levels $\alpha_p = 0.45, 0.9$ of uncertainty in SNR arising from uncertain drift rate. Nominal MMPC for $\eta = \tilde{\eta}$ (solid), and extreme MMPCs for $\eta = \tilde{\eta}(1 + \alpha_p)$ (dashed) and for $\eta = \tilde{\eta}(1 - \alpha_p)$ (dot-dashed); the latter coincides with the OPC for all α_p . Performance bands for SNR within the range consistent with the presumed uncertainty are bounded by the extreme MMPC and the OPC. Circles on the OPC and nominal MMPCs mark the corresponding performance parameters under the same conditions ($\tilde{\eta} = \eta = 2$ and $D_{tot} = D = 1$)

If the SNR varies within the range consistent with the presumed level of uncertainty, i.e., $\eta \in [\tilde{\eta}(1 - \alpha_p), \tilde{\eta}(1 + \alpha_p)]$ for $\alpha_p \leq 1$, the maximin performance bands are bounded by the OPC (Eq. (8), obtained for $\eta = \tilde{\eta}(1 - \alpha_p)$) and by Eq. (33) with $\chi = \frac{1 - \alpha_p}{1 + \alpha_p}$ (for $\eta = \tilde{\eta}(1 + \alpha_p)$). Fig. 5 shows MMPCs and performance bands for two levels of presumed uncertainty. The curves exhibit a strong leftward shift in peak values, and lie below the OPC for most error rates, deviating even further from the experimental data than the MMPCs shown in Fig. 4, and are therefore excluded from the comparisons performed in §4.

We stress that even if an MMPC lies below the OPC, implying shorter mean reaction times for the same error rates, it nevertheless results in poorer performance! Performance curves cannot be compared directly for a given error rate (or reaction time), but should be compared for the same operating conditions, as depicted by the three circles in Fig. 5. Each circle identifies the performance parameters associated

with the operating curve on which it resides when $\tilde{\eta} = \eta = 2$ and $D_{tot} = D = 1$. The resulting reward rate decreases from 0.707 for the OPC, to 0.706 for the MMPC with $\alpha_p = 0.45$ and 0.665 for the MMPC with $\alpha_p = 0.9$. As the presumed level of uncertainty increases, performance under the maximin strategy degrades, but that performance level can be assured under wider conditions (as specified by the Info-gap and the presumed level of uncertainty).

Robust satisficing and opportunity facilitating strategies: Under uncertainties in the SNR that arise from uncertainties in the drift rate, the robust satisficing and opportunity facilitating strategies result in performance curves that further deviate from the experimental data, as described in Appendix B.8. These curves lie below the OPC for intermediate error rates and reach plateaus for near-chance error rates, as shown in Fig. 11. These strategies therefore cannot explain the observed SATs and are not considered in the detailed data comparison in §4.

3.2.3 Trial-to-trial variations in signal-to-noise ratio

As noted above, trial-to-trial variations in the SNR are equivalent to trial-to-trial variations in the drift rate A with fixed noise variance σ^2 , which has been advanced to account for differences between error and correct mean response times [3, 4, 5]. Here we extend the analysis of §3.2.2 to include uncertainties that may arise from both estimation errors and trial-to-trial variations in the SNR, under fixed noise variance. However, unlike [3, 4, 5], where a Gaussian distribution of drift rates is used, here we assume a series of bounded SNRs:

$$U_{ser}(\alpha, \tilde{\eta}) = \{ \{ \eta_i \}_{i=1}^{\infty} : \eta_i > 0, |\eta_i - \tilde{\eta}| \leq \alpha \tilde{\eta} \}, \quad (34)$$

where $\sqrt{\eta_i}$ denotes the SNR on the i^{th} trial.

The robust analysis is distribution-free and based on worst case scenarios. Under fixed noise variance σ^2 , the worst RR with any series of SNRs within the Info-gap model of (34), is obtained when the SNR is fixed at the worst (lowest) level consistent with the level of uncertainty, as stated and proved in the next Theorem:

Theorem 3.5. *Worst performance with a variable SNR: Consider a 2AFC task with fixed threshold-to-noise ratio $\vartheta = x_{th}/\sigma$. The worst performance that can occur with any series of bounded SNRs consistent with the Info-gap model of Eq. (34) coincides with the worst performance that can occur with a fixed SNR consistent with the Info-gap model of Eq. (26) at the same level of uncertainty.*

Proof: See Appendix B.9. □

Theorem 3.5 implies that the performance curves for the variable SNR case considered here are the same as for the fixed SNR case considered in §3.2.2. It was concluded there that under uncertainties in the drift rate neither the MMPCs nor the RSPCs and OFPCs match the experimental data as well as the MMPCs for uncertainty in noise variance. Thus, robust strategies for uncertainties in drift rate cannot account for the observed SATs, even when they include bounded trial-to-trial variations.

3.3 Summary of robust strategies

We have derived two classes of performance curves, MMPCs and RSPCs, to describe robust behaviors in response to uncertainties in intertrial delays and in SNR, the latter deriving from either noise variance or drift rate. The general forms of the resulting curves appear to qualitatively match the data of Fig. 1 only in the three cases of MMPCs and RSPCs for delay uncertainty (Figs. 2 and 3), and MMPCs for SNR uncertainty due to uncertain noise variance (Fig. 4), and so we include only these three cases in the comparisons of §4. However, for completeness and for use with other data sets, examples of the remaining performance curves are given in Figs. 5 above, and in Figs. 10 and 11 in Appendix B.

4 Comparisons with experimental data

We now further evaluate the three robust performance curves that qualitatively match the behavioral data described in §2.2, and compare how they, and the accuracy weighted curves of §2.3, fit this data.

Experimental data: We limit our comparison to data from blocks with no penalty delays $D_p = 0$ i.e., $D_{tot} = D$), because OPCs for RR_m and RA are only defined for $D_p = 0$, and RSPCs for uncertainty in D involve the ratio D/D_{tot} , which differs between blocks with $D_p = 0$ and $D_p > 0$.

As noted in §2, the speed-accuracy trade-off differs between subjects with different total earnings, and hence performance curves must be compared with data from different groups of subjects separately. In particular, some subjects with the lowest earnings had decision times an order of magnitude higher than others (see the bottom panel of Fig. 6), and we therefore analyse the lowest 10% of subjects separately. We then split the remaining 90% of subjects into 3 equal groups, but upon finding no significant difference between mean normalized decision times of the 30 – 60% group and the 60 – 90% group (paired t-test across 10 ranges of error rates: $p > 0.35$), we henceforth pooled the data from these groups. The normalized decision times for the resulting middle 60% and for the top 30% groups are shown in the upper panels of Fig. 6. The number of experimental conditions falling into each bin (i.e. the number of data points averaged to obtain each mean decision time) is given in Table 1.

Fitting method: Except the OPC for RR , which is parameter-free, each performance curve includes one free parameter, denoted by q or α . We estimated these parameters via the maximum likelihood method, using the standard assumption that normalized decision times are normally distributed around the performance curve being fitted, with variance estimated from the data. Further comments appear below and in Appendix C.

Three of the performance curves are defined only for a restricted range of $p(\text{err})$ s, but subjects exhibit $p(\text{err})$ s from the entire range $[0, 0.5]$ (see Table 1). We addressed the issue of undefined values in the following ways. Firstly, RSPCs were extended to include the OFPC when q values corresponded to a level of windfall performance,

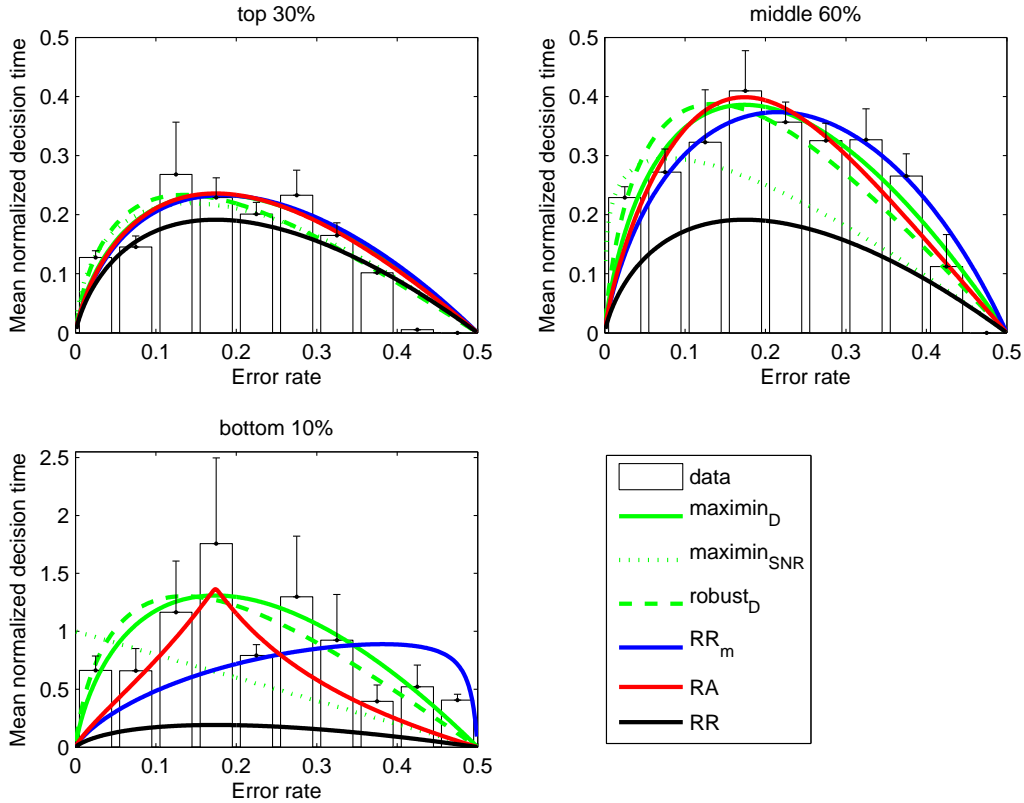


Figure 6: Comparisons of performance curves with mean normalized decision times from experimental blocks with $D_p = 0$ for three groups of subjects sorted by total rewards acquired. Error bars show standard error of mean normalized decision times. Note the difference in scale of the vertical axes in upper and lower panels. Here and throughout this section, $\text{maximin}_{\text{SNR}}$ refers to the MMPC due to uncertainty in noise variance.

as described in §3.1, Fig 3. This corresponds to the assumption that when subjects cannot guarantee the required level of performance, even under nominal conditions, they attempt to facilitate the opportunity for achieving it. Secondly, the OPC for RA is not defined around the minimum of \mathcal{E} for $q > 1.096$ (cf. inequality (12)), so in fitting it we restricted the weight parameter to $q \leq 1.096$. Thirdly, the OPC for RR_m is not defined for $p(\text{err})$ s close to 0.5 if $q > 1$, so in fitting it we restricted the weight parameter to $q \leq 1$.

Results: Fig. 6 shows the best-fitting performance curves for the three groups, and parameter values estimated for each performance curve are given in Table 2. For each performance measure, the corresponding parameter α or q increases as overall performance decreases, being lowest for the top 30% and highest for the bottom 10% of subjects. This is consistent with natural interpretations of the various theories:

Error rate range	Top 30%	Middle 60%	Bottom 10%
0 – 5%	37	82	9
5 – 10%	17	32	7
10 – 15%	12	12	3
15 – 20%	12	23	6
20 – 25%	19	35	2
25 – 30%	15	39	5
30 – 35%	7	13	4
35 – 40%	1	6	2
40 – 45%	1	5	2
45 – 50%	4	3	2

Table 1: Numbers of experimental conditions falling into each error rate bin for the three panels of Fig. 6.

(i) according to MMPCs it implies that poorer performers have higher presumed uncertainty in D or SNR; (ii) according to RSPCs it implies that they require lower levels of reward, and (iii) according to the RR_m and RA criteria it implies that they place higher emphasis on accuracy. Note that the best fit for the RA theory with the lowest 10% is with $q = 1.096$, at the constraint limit of Eq. (12), and that the corresponding OPC has become concave on either side of its sharp peak.

Performance curve	Top 30%	Middle 60%	Bottom 10%
MMPC for D	0.22	1.02	5.84
MMPC for SNR	0.19	0.54	∞
RSPC for D	0.72	1.85	8.57
OPC for RR_m	0.14	0.49	0.98
OPC for RA	0.15	0.55	1.096

Table 2: Values of performance curve parameters (q or α) estimated using the maximum likelihood method from the data from the three groups of subjects sorted by total earning. For the MMPC with uncertainty in SNR and bottom 10%, the higher the value of α the higher the likelihood of the data. Key: MMPC: maximin performance curve; RSPC: robust satisficing performance curve; OPC: optimal performance curve; SNR: signal-to-noise ratio due to noise variance; D : delays; RR_m : modified reward-rate Eq. (13); RA : reward-accuracy rate Eq. (9).

Comparison of likelihoods of data given different performance curves reveals that the MMPC with uncertainty in D fits the data best for all three groups of subjects. Table 3 shows the ratios of the likelihood of the data given MMPC with uncertainty in D , with likelihoods given the other performance curves. For the top 30% of subjects all curves fit comparably well, except the OPC for RR . For the middle and bottom groups the differences in fit qualities increase significantly, especially in comparison to the OPCs for RR and RR_m , and the MMPC for uncertainty in SNR.

The last column of Table 3 shows likelihood ratios of the data from all subjects, given the different performance curves. It indicates that the data is several orders of

Performance curve	Top 30%	Middle 60%	Bottom 10%	Product
MMPC for SNR	1.13	$> 10^8$	308.38	$> 10^{10}$
RSPC for D	1.99	4.04	1.72	13.82
OPC for RR_m	1.28	38.79	208.47	$> 10^4$
OPC for RA	1.11	8.11	4.83	43.48
OPC for RR	47.34 (17.42)	$> 10^{21}$	$> 10^6$	$> 10^{28}$

Table 3: Likelihood ratios of the data given the MMPC with uncertainty in D compared to other performance curves (rows), for different groups of subjects (columns). The last column is the product of the first three and thus expresses likelihood ratios for all data. The higher the likelihood value, the less likely is it that the data could be generated by the corresponding model, in comparison to the MMPC with uncertainty in D . To account for the fact that all curves have one free parameter except the OPC, the bottom row includes in brackets the ratio of the likelihood of the data given the two curves divided by $\exp(1)$, as suggested by Akaike [49]: see Appendix C. Key: MMPC: maximin performance curve; RSPC: robust satisficing performance curve; OPC: optimal performance curve; SNR: signal-to-noise ratio due to noise variance; D : delays; RR_m : modified reward-rate Eq. (13); RA : reward-accuracy rate Eq. (9).

magnitude more likely given MMPC with uncertainty in D than given OPC for either RR or RR_m , or for the MMPC with uncertainty in SNR. It also shows that the data is over 13 times more likely given the MMPC with uncertainty in D than the RSPC with uncertainty in D , and over 43 times more likely given MMPC with uncertainty in D than given OPC for RA . In summary, the MMPC with uncertainty in D fits the experimental data for all three subgroups better than any other performance curve, although the differences in fit quality among this one, the RSPC for D , and the OPC for RA are relatively small. Additional details are given in Appendix C.

At least two effects influence the reliability of the likelihood ratio estimates. First, the test assumes that experimental decision times are normally distributed around the performance curve being fitted. We assessed this as described in Appendix C, concluding that there is no evidence for non-Gaussianity in the majority of bins, although up to 3 out of the 10 distributions in each group are significantly non-Gaussian, being skewed toward long DTs. Second, likelihood ratios may depend on the way that subjects are split into groups, since we implicitly assume a homogeneous strategy within each group. To investigate this we fitted data from six additional splits, ranging from all 80 individuals treated singly, to two groups of 40 subjects each. As described in Appendix C, while different splits lead to different likelihood ratios, the relative ordering of performance curves remains essentially the same as that of Table 3. We therefore believe that the conclusions drawn above and developed in §6, are sound.

5 Comparison of optimal and robust strategies

The Info-gap approach addresses the effects of discrepancies between estimated (nominal) parameters and actual ones. Here we compare the influence of such discrepancies on performance curves found by optimizing RR directly, but for incorrect parameter values, with their effects on MMPCs obtained by the robust techniques of §3.

Suboptimal performance curves: The derivation of the OPC of §2 presumes that subjects estimate SNRs and delays correctly. Maximizing RR for incorrect parameters results in *sub-optimal performance curves* (sub-OPCs) that would yield maximum RR s for the specific mis-estimated parameters instead of the actual ones.

To facilitate comparison with the robust performance curves of §3 we express potential discrepancies in estimated delays and SNR as $\tilde{D}_{tot} = D_{tot}/(1 \pm q_D)$, and $\tilde{\eta} = \eta/(1 \pm q_\eta)$, respectively. Assuming discrepancies only in delays, the resulting sub-OPCs are parameterized by q_D and have the same form as MMPCs for uncertainties in delays, Eq. (20), with $\gamma = (1 \pm q_D)^{-1}$. Similarly, for discrepancies in SNR alone, the sub-OPCs are parameterized by q_η and have the same form as the MMPCs for uncertainties in SNR, Eq. (30), with $\chi = (1 \pm q_\eta)^{-1}$. These two families of sub-OPCs are shown in Figs. 7 and 8. The sub-OPC with discrepancies in delays provides a better match to the overall data, capturing the larger $\langle DT \rangle$ range for $p(\text{err}) \approx 20\%$.

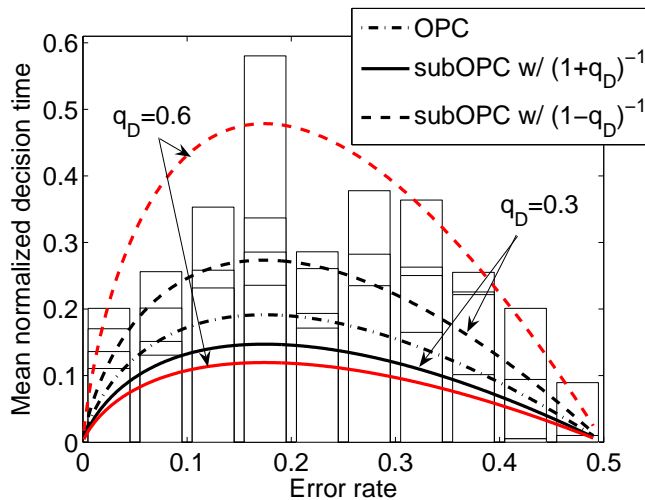


Figure 7: Optimal ($q_D = 0$) and sub-optimal performance curves (subOPCs) in which threshold is optimized with respect to estimated total delays D_{tot} differing from the actual value by $1/(1 + q_D)$ (solid) and $1/(1 - q_D)$ (dashed).

Comparison of sub-OPCs and MMPC: Selecting the threshold that maximizes the worst case RR in the presence of uncertainties in delays results in MMPCs that agree well with the experimental data (Fig. 6). The sub-OPCs of Fig. 7, generated when performance is optimized for poorly estimated delays, are also capable of




Figure 8: Optimal ($q_\eta = 0$) and sub-optimal performance curves (subOPCs) in which threshold is optimized with respect to estimated SNRs $\sqrt{\eta}$, where η differs from the actual value by $1/(1 + q_\eta)$ (solid) and $1/(1 - q_\eta)$ (dashed).

matching the data. However, there are subtle differences between MMPCs and sub-OPCs, which may permit one to distinguish which strategy subjects might favor.

The most significant difference relates to conditions that would generate normalized decision times *below* the OPC for given $p(\text{err})$. Under the maximin strategy, normalized decision times always lie above the OPC provided that variations in delays occupy the Info-gap set defined by the presumed uncertainty level α_p (recall Fig. 2, in which the OPC forms the lower bound for all performance bands). Only when delays exceed the upper bound of the Info-gap model can normalized decision times fall below the OPC (see Eq. (20) and the discussion following it). In contrast, when performance is optimized, it is sufficient for delays to be underestimated to result in points below the OPC, as shown by the solid curves for $1/(1 + q_D)$ in Fig. 7. Hence, while normalized decision times might lie below the OPC under both strategies, they are less likely to do so under the maximin strategy. Data points below the OPC are scarce and thus better support the maximin strategy: see Fig. 6.

Sub-OPCs and MMPCs also require different levels of discrepancy or uncertainty to generate the same upper bound for the performance band. When delays vary within a range consistent with presumed uncertainty α_p , Eq. (20) with $\gamma = (1 + \alpha_p)/(1 - \alpha_p)$ gives this upper bound, as the discussion following that equation indicates. In contrast, the upper sub-OPC is given by Eq. (20) with $\gamma = (1 - q_D)^{-1}$. So for $\alpha_p < 1$, the same upper bound is achieved with $q_D > \alpha_p$: for example, the extreme MMPC for $\alpha_p = 0.4$ of Fig. 2 is similar to the sub-OPC with $q_D = 0.6$ of Fig. 7. Noting that equal values of α_p and q_D describe the same discrepancy between estimated and actual delays, we conclude that a smaller level of uncertainty can explain the experimental results under the assumption that the maximin strategy is used, than under the assumption that the optimal strategy is (mis)used.

6 Discussion

We have developed an approach to two-alternative forced choice (2AFC) tasks that accounts for potential uncertainties in experimental conditions, specifically, estimated delays and signal-to-noise ratio. Two strategies – maximin and robust satisficing – were proposed. The former assumes that subjects maximize guaranteed performance under a presumed uncertainty level; the latter assumes that they maximize the uncertainty under which a required level of performance can be guaranteed. We compared these strategies with optimal procedures based on reward rate and modified reward rates weighted for accuracy, and found that the maximin strategy with uncertainties in delays predicts performance curves that best match behavioral data from a group of 80 subjects (Fig. 6). Performance curves predicted by the robust satisficing strategy with uncertainties in delays, and optimization of a modified reward-accuracy rate (RA) are the closest competitors. Additional experiments are needed to further assess which strategy is more consistent with human decision making, as further discussed below.

The maximin and robust satisficing strategies (for uncertainties in the delays) result in performance bands around the nominal performance curves. These bands explain well the range of decision times ($\langle DT \rangle / D_{tot}$) observed for given error rates. We also evaluated the effect of incorrectly-estimated delays on the optimal strategy and derived the resulting performance bands (Fig. 7). While these also provide good descriptions of the experimental results, subtle differences suggest that the maximin strategy is more consistent with behavior in the present experiments; but again, additional experiments are needed to be conclusive.

Our Info-gap framework allows uncertainty in the parameters describing the pure DDM, and in doing so provides an alternative to the probabilistically-extended DDM that includes trial-to-trial parameter variations [5]. We show that uncertainty in SNR can derive from trial-to-trial variations in drift rate, albeit bounded ones, and thus can also account for differences between response times for correct and error trials. However, unlike those for uncertainties in delays, we find that the resulting performance curves do not match well the experimental data (nor do performance curves for Gaussian-distributed drift rates [20]). This does not mean that trial-to-trial variations in SNR are absent, but it may imply that the presumed uncertainty in the SNR is much smaller than the presumed uncertainty in delays.

It is notable that uncertainties in delays result in performance curves that fit the data well, while uncertainties in SNR do not do so. A major source of the former may be attributed to the scalar property of interval timing [47, 48]. Psychophysical experiments suggest that elapsed time estimates are distributed normally around the actual value with a standard deviation proportional to it. Thus the sizes of confidence intervals around an estimated experiential delay are proportional to the delay itself. This relationship is captured by the proportional Info-gap model of §3.1. In contrast to potential uncertainties in estimating time intervals, human subjects seem to be very sensitive to SNR levels. For example, Luijendijk [50] found that the perception threshold for noise in images was as much as 27dB below signal level. This suggests that humans can accurately assess SNR, although further, more direct,

experimental verification would be required.

The underlying assumption in this paper is that subjects either optimize or satisfice by setting their decision thresholds based on estimates of intertrial delays and SNR, as modeled by a DD process. In a related study, Simen et al. [37] proposed a neural network model for rapid threshold adjustment based on estimates of reward rate RR that are updated as trials proceed. Their model relies on prior, long-term, learning of an approximate linear relationship between maximum RR and threshold that is independent of SNR over an appropriate SNR range. It too implicitly assumes the ability to estimate time intervals, since the discrete rewards following correct responses must be converted into RR , and it also suggests that SNR may not play as important a role as delays.

We propose additional experiments to investigate whether subjects satisfice or optimize in 2AFC, along two lines of research: (i) direct investigation of the effects of induced uncertainties, and (ii) correlation of performance in 2AFC with accuracy in interval timing. Uncertainties could be induced by varying SNRs or delays randomly in each block of trials, although the former violates the assumption of statistical stationarity on which the optimality proofs for the SPRT and DD model rely. Variation of response-to-stimulus interval D_{RSI} about a fixed mean is possible, but a better approach might be to keep delays and SNR fixed within relatively short blocks, and draw from narrow or broad distributions from block to block. The resulting effects on performance would then be assessed in comparison with the different decision strategies. Alternatively, a subject's standard deviation in interval timing, σ_I , could be measured directly in separate time-estimation trials, and then correlated with performance on 2AFC tasks. The maximin strategy suggests that the normalized decision time $\langle DT \rangle / D_{tot}$ is proportional to $\gamma = 1 + \alpha_p$. Assuming that the presumed uncertainty α_p is proportional to σ_I , there should be a correlation (across subjects) between $\langle DT \rangle / D_{tot}$'s in 2AFC experiments and average errors in interval timing. One could also correlate timing ability with performance on a deadlined choice task with substantial penalties for failures to respond before the deadline. Unlike the free response paradigm, poor timers are predicted to respond prematurely, and hence *faster and less accurately*, under such conditions [51]: in the opposite direction to their suboptimal slower and more accurate free response behavior.

The robust satisficing strategy is relevant when a critical performance level must be met, while the maximin strategy is appropriate for optimizing performance under an expected level of uncertainty. The current results indicate that the majority of subjects in 2AFC tasks seem to follow maximin rather than robust satisficing or optimal performance curves. It would also be interesting to investigate whether: (i) performance curves for individuals follow the maximin curves of Eq. (20) (Fig. 2) or the corresponding robust satisficing curves of Eq. (24) (Fig. 3), and (ii) if subjects change their strategy under different conditions, i.e., do they resort to robust satisficing when the success of the whole block (and the consequent net reward) depends on achieving a specific level of performance. To address the first question, large data sets with broad SNR variation would be needed for each subject, to cover the entire $p(\text{err})$ range. The second question demands additional experiments in which

a fixed reward is given when outperforming a preset required level, instead of the current strategy in which the reward increases linearly with the number of correct responses.

As noted in §1, the pure DD model is analytically tractable, yielding explicit, parameter-free OPCs. It supports a normative decision theory that can explain sub-optimal experimental data by allowing analytical derivation of one-parameter models for comparison with the optimal strategy that maximizes reward rate. The current analysis suggests that robust strategies against uncertainties in delays best explain the sub-optimal results in the present data. Nevertheless, additional potential sources of nonoptimality identified in §2.1 should also be investigated.

Appendices: Proofs and mathematical details

A Reward rate and optimal performance

A.1 Reward rate

Lemma A.1. Let the random variable S_t denote the number of correct decisions in time t . As $t \rightarrow \infty$, the mean number of correct decisions per unit time $\langle S_t \rangle / t$ approaches the RR defined by Eq. (4).

Proof: Consider the DD model of Eq. (1), which is re-initialized immediately after crossing the threshold. The resulting passage times define a sequence of points in time, and hence may be regarded as a realization of a point process. Furthermore, since the first passage time is independent of the previous passage times, the resulting point process can be described as a renewal process [52], i.e., a point process in which intervals between the points are the realizations of a random variable T whose pdf depends only on the time since the last point (i.e., last passage time). Let N_t be the random variable indicating the number of passages in time t . For a renewal process the asymptotic distribution of N_t for large t is normal with mean $\langle N_t \rangle = t / \langle T \rangle$ [52, Eq. (3.3.3)] (this relies on the asymptotic normal distribution of the sum of the first r renewals for large r). So, for large t , $t = \langle T \rangle \langle N_t \rangle$.

Let $p(\text{corr})$ be the probability that a decision is correct. Then, for n decisions the mean number of correct decisions is: $\langle S_t | N_t = n \rangle = p(\text{corr}) n$; and the mean number of correct decisions is:

$$\langle S_t \rangle = \sum_{n=0}^{\infty} \langle S_t | N_t = n \rangle \Pr[N_t = n] = p(\text{corr}) \sum_{n=0}^{\infty} n \Pr[N_t = n] = p(\text{corr}) \langle N_t \rangle. \quad (35)$$

Hence, asymptotically as $t \rightarrow \infty$ we have

$$\frac{\langle S_t \rangle}{t} = \frac{p(\text{corr}) \langle N_t \rangle}{\langle T \rangle \langle N_t \rangle} = \frac{p(\text{corr})}{\langle T \rangle}. \quad (36)$$

The proof is completed by noting that the decision times form a renewal process too, where each decision interval equals the corresponding first passage time plus

processing and delay intervals. Substituting $p(\text{corr}) = 1 - p(\text{err})$ and $\langle T \rangle = \langle DT \rangle + T_0 + D_{RSI} + D_{pen}p(\text{err})$ in Eq. (36) results in Eq. (4). \square

A.2 Unique Maximum of the OPC

Here we show that the function \mathcal{E} defined in Eq. (11) has a unique (global) minimum and estimate its value. This allows us to state explicit conditions on the weight parameter q that guarantee that the nondimensional decision times given by Eq. (10) are real.

Lemma A.2. The function \mathcal{E} of Eq. (11) has a unique minimum $\mathcal{E} \approx 5.224$ at $ER \approx 0.1741$ and \mathcal{E} approaches $+\infty$ as $ER \rightarrow 0$ and as $ER \rightarrow 0.5$.

Proof: The limits at $ER = 0$ and 0.5 are obtained directly for the second term in \mathcal{E} , and via L'Hôpital's rule for the term involving logarithms. Before computing the derivative to prove that there is a unique minimum, it is convenient to change variables by defining

$$y = \frac{1 - p(\text{err})}{p}(\text{err}) \Rightarrow p(\text{err}) = \frac{1}{1 + y}; \quad (37)$$

note that as $p(\text{err})$ rises from 0 to 0.5, y falls monotonically from $+\infty$ to 1. We then have

$$\frac{d\mathcal{E}}{dy} = \frac{d}{dy} \left[\frac{y+1}{\log(y)} + \frac{y+1}{y-1} \right] = \frac{\log(y) - 1 - y^{-1}}{[\log(y)]^2} - \frac{2}{(y-1)^2}. \quad (38)$$

Setting Eq. (38) equal to zero and rearranging, we find that critical points of \mathcal{E} occur at the roots of

$$\log(y) - y^{-1} = 1 + 2 \left[\frac{\log(y)}{y-1} \right]^2, \quad (39)$$

but in fact the solution is unique, since the left hand side of (39) is strictly increasing from -1 to ∞ and its right hand side monotonically decreases from 2 to 1, as y goes from 1 to ∞ . The former claim is easy to check, and the latter follows from computation of the derivative of $[\log(y)/(y-1)]$:

$$\frac{d}{dy} \left[\frac{\log(y)}{y-1} \right] = \frac{1 - y^{-1} - \log(y)}{(y-1)^2}. \quad (40)$$

This expression is clearly negative for all $y \geq e$, and the following series expansion shows that it is in fact negative for all $y > 1$ and zero only at $y = 1$:

$$\log(y) = \sum_{j=1}^{\infty} \frac{1}{j} \left(\frac{y-1}{y} \right)^j = 1 - y^{-1} + \sum_{j=2}^{\infty} \frac{1}{j} \left(\frac{y-1}{y} \right)^j, \quad (41)$$

see [53, Formula 4.1.25]. Numerical solution of Eq. (39) yields the estimates of the lemma. \square

B Robust performance

B.1 Performance under uncertainties in delays

To facilitate computation we express performance in terms of inverse reward rate $IR = RR^{-1}$ instead of reward rate RR . Following Eqs. (4) and (5) (see also Eq. (7)), IR can be expressed as a function of $p(\text{err})$ and $\langle DT \rangle$:

$$IR(p(\text{err}), \langle DT \rangle : D, D_{tot}) = \frac{\langle DT \rangle + D + (D_{tot} - D)p(\text{err})}{1 - p(\text{err})}. \quad (42)$$

and as a function of θ and η :

$$IR(\theta, \eta : D, D_{tot}) = \theta(1 - \exp(-2\eta\theta)) + D + D_{tot} \exp(-2\eta\theta). \quad (43)$$

The worst performance (maximum IR) that may occur under uncertainty α in the Info-gap model of Eq. (15) is denoted by: $I_{MAX}(\alpha : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) \equiv \left(\max_{D, D_{tot} \in U(\alpha, \tilde{D}, \tilde{D}_{tot})} IR \right)$.

According to Eq. (43), IR is maximized with respect to delays when $D = \tilde{D}(1 + \alpha)$ and $D_{tot} = \tilde{D}_{tot}(1 + \alpha)$, so:

$$\begin{aligned} I_{MAX}(\alpha : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) &= \theta + \tilde{D}(1 + \alpha) + [\tilde{D}_{tot}(1 + \alpha) - \theta] \exp(-2\eta\theta) \\ &= I_0(\theta, \eta) + (1 + \alpha)I_D(\theta, \eta : \tilde{D}, \tilde{D}_{tot}); \end{aligned} \quad (44)$$

where $I_0(\theta, \eta) \equiv \theta(1 - \exp(-2\eta\theta))$ is the inverse rate with zero delays, and $I_D(\theta, \eta : D, D_{tot}) \equiv D + D_{tot} \exp(-2\eta\theta)$ is the additional inverse rate due to the delays.

Similarly, the best performance (minimum IR) that can be obtained is denoted as: $I_{MIN}(\alpha : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) \equiv \left(\min_{D, D_{tot} \in U(\alpha, \tilde{D}, \tilde{D}_{tot})} IR \right)$. According to Eq. (43),

IR is minimized with respect to delays when $D = \max(0, \tilde{D}(1 - \alpha))$ and $D_{tot} = \max(0, \tilde{D}_{tot}(1 - \alpha))$. For $\alpha \leq 1$, Eq. (43), implies that:

$$\begin{aligned} I_{MIN}(\alpha \leq 1 : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) &= \theta + \tilde{D}(1 - \alpha) + [\tilde{D}_{tot}(1 - \alpha) - \theta] \exp(-2\eta\theta) \\ &= I_0(\theta, \eta) + (1 - \alpha)I_D(\theta, \eta : \tilde{D}, \tilde{D}_{tot}). \end{aligned} \quad (45)$$

The case $\alpha \geq 1$ includes the favorable condition of zero delays, so $I_{MIN}(\alpha \geq 1 : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) = I_0(\theta, \eta)$.

B.2 Robustness curves for uncertainties in delays

Here we derive and depict the robustness curves for uncertainties in delays.

Theorem B.1. Robustness curves for 2AFC under uncertainties in the delays: *Given the SNR and nominal delays, the robustness depends on the required performance R_r as:*

$$\hat{\alpha}(R_r : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) = \max \left\{ 0, \frac{1/R_r - \theta(1 - \exp(-2\eta\theta))}{\tilde{D} + \tilde{D}_{tot} \exp(-2\eta\theta)} - 1 \right\} \quad (46)$$

Proof: The internal minimization of Eq. (21) is derived in Appendix B.1, and expressed in Eq. (44) as the maximum inverse reward rate. Accordingly, the internal inequality can be expressed as:

$$\theta + \tilde{D}(1 + \alpha) + [\tilde{D}_{tot}(1 + \alpha) - \theta] \exp(-2\eta\theta) \leq R_r^{-1} \quad (47)$$

The robustness is obtained when equality is achieved, so Eq. (47) implies that the robustness is given by Eq. (46). \square

Eqs. (5) and (46) imply that the nominal performance that can be achieved with the nominal delays, $RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot})$, has zero robustness (see also Fig 9).

According to Eq. (46), the robustness curves depends linearly on the inverse of the required performance $I_r = 1/R_r$, as depicted in Fig. 9 (right panel) for the specific case $\tilde{D}_{tot} = \tilde{D} = 1$, and $\eta = 1$. The optimal threshold ($\theta_{op} \approx 0.3960$), achieves optimal performance (minimal IR , here $IR_{op} \approx 1.670$), with zero robustness. Only sub-optimal performance levels can be guaranteed robustly. Furthermore, the robustness curve for $\theta > \theta_{op}$ crosses the robustness curve for θ_{op} . Thus, when a sub-optimal performance is required, it can be achieved with higher robustness to uncertainties using a sub-optimal threshold than the optimal threshold.

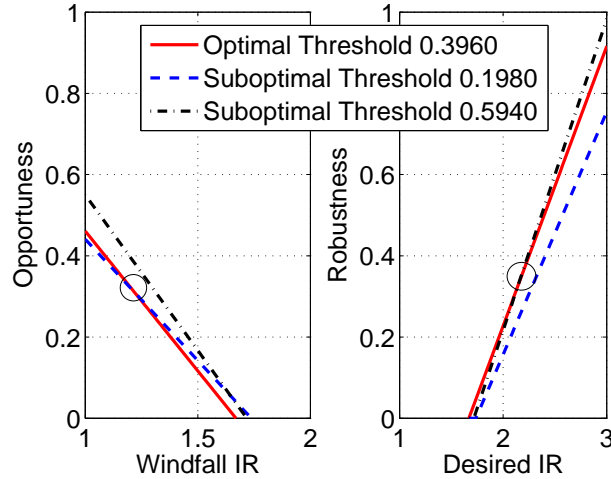


Figure 9: Robustness curves (right) and opportuneness curves (left) under uncertainties in delays for the optimal threshold, a lower and a higher thresholds, exhibiting curve crossing as indicated by the circles. Here $\theta_{op} \approx 0.3960$ for $\tilde{D}_{tot} = \tilde{D} = 1$ $\eta = 1$. See text for discussion.

B.3 Robust satisficing performance curves under uncertainties in delays

Here we prove Theorem 3.2, which states that under uncertainties in the delays, θ_{RS} satisfies Eq. (23). We also show that the constraint on the required performance R_r implies that the RSPCs lie above the OPC.

For a fixed level of required performance $R_r < RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot})$, the derivative of $\hat{\alpha}(R_r : \theta, \eta, \tilde{D}, \tilde{D}_{tot})$ (Eq. 46) with respect to θ can be expressed as

$$\frac{\partial \hat{\alpha}}{\partial \theta} = \frac{g(\theta, \eta : R_r, \tilde{D}, \tilde{D}_{tot}) \exp(-2\eta\theta)}{\left(\tilde{D} + \tilde{D}_{tot} \exp(-2\eta\theta)\right)^2}, \quad \text{where} \quad (48)$$

$$g(\theta, \eta : R_r, \tilde{D}, \tilde{D}_{tot}) = \frac{2\eta\tilde{D}_{tot}}{R_r} - [2\eta\theta + \exp(2\eta\theta) - 1] \tilde{D} - [2\eta\theta - \exp(-2\eta\theta) + 1] \tilde{D}_{tot}. \quad (49)$$

Hence, extrema of $\hat{\alpha}(R_r : \theta, \eta, \tilde{D}, \tilde{D}_{tot})$ as a function of θ satisfy $g(\theta, \eta : R_r, \tilde{D}, \tilde{D}_{tot}) = 0$. Since the second derivative is always negative,

$$\begin{aligned} \left. \frac{\partial^2 \hat{\alpha}}{\partial \theta^2} \right|_{\frac{d\hat{\alpha}}{d\theta}=0} &= \frac{\exp(-2\eta\theta) \partial g(\theta, \eta : R_r, \tilde{D}, \tilde{D}_{tot}) / \partial \theta}{\left(\tilde{D} + \tilde{D}_{tot} \exp(-2\eta\theta)\right)^2} \\ &= \frac{-2\eta \exp(-2\eta\theta) \left[\tilde{D}(1 + \exp(2\eta\theta)) + \tilde{D}_{tot}(1 + \exp(-2\eta\theta)) \right]}{\left(\tilde{D} + \tilde{D}_{tot} \exp(-2\eta\theta)\right)^2} < 0, \quad (50) \end{aligned}$$

these extrema are maxima. The condition $g(\theta, \eta : R_r, \tilde{D}, \tilde{D}_{tot}) = 0$ defines a unique maximum and, from Eq. (49), the corresponding RS threshold obeys Eq. (23). \square

The robust satisficing strategy is relevant only when $R_r < RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot})$. According to Eq. (43), this constraint implies that the RSPC holds only when $(R_r^{-1} - \tilde{D})/\tilde{D}_{tot} \geq (1 - p(\text{err}))^{-1}(p(\text{err}) + \langle DT \rangle / \tilde{D}_{tot})$, or alternatively when

$$\langle DT \rangle / \tilde{D}_{tot} \leq q - p(\text{err})(1 + q), \quad (51)$$

where $q \equiv (R_r^{-1} - \tilde{D})/\tilde{D}_{tot}$. The maximum value of $RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot})$ is obtained with the optimal threshold defined by Eq. (6), or equivalently at the OPC defined by Eq. (8). Thus, the constraint $R_r \leq RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot})$ forces RSPCs to lie above the OPC, which is obtained when $R_r = RR(\theta_{op}, \eta : \tilde{D}, \tilde{D}_{tot})$.

B.4 Opportunity facilitating performance curves under uncertainties in delays

A better than optimal performance level is considered a *windfall performance level*. While windfall performance cannot be guaranteed under nominal conditions even with the optimal strategy, it may occur under favorable conditions consistent with the uncertainty level.

Definition B.1. Opportuneness with uncertainties in delays: *Consider a 2AFC strategy with threshold θ and a windfall performance R_w . The opportuneness $\hat{\beta}$ for gaining R_w under uncertainties in delays is the lowest horizon of uncertainty in the Info-gap model of Eq. (15) that provides the opportunity for gaining R_w [23]. The*

opportuneness is infinite if R_w can never occur under any level of uncertainty in the delays:

$$\hat{\beta}(R_w : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) = \min \left\{ \alpha \left| \left(\max_{D, D_{tot} \in U(\alpha, \tilde{D}, \tilde{D}_{tot})} RR \right) \geq R_w \vee \alpha = \infty \right. \right\}. \quad (52)$$

The functional dependence of the opportuneness on the windfall performance is termed the opportuneness curve.

Theorem B.2. Opportuneness curves for 2AFC under uncertainties in the delays: Given the SNR and nominal delays, the opportuneness depends on the windfall performance R_w as:

$$\hat{\beta}(R_w : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) = \min \left\{ 1, \max \left(0, 1 - \frac{R_w^{-1} - \theta(1 - \exp(-2\eta\theta))}{\tilde{D} + \tilde{D}_{tot} \exp(-2\eta\theta)} \right) \right\} \quad (53)$$

Proof: The maximum RR in the internal maximization of Eq. (52) is the inverse of the minimum IR given by Eq. (45). Hence, the internal inequality can be expressed as:

$$I_{MIN}(\alpha : \theta, \eta, \tilde{D}, \tilde{D}_{tot}) \leq R_w^{-1} \quad (54)$$

The opportuneness is obtained when equality is achieved, so Eq. (45) implies that the opportuneness is given by Eq. (53). \square

Examples of opportuneness curves are depicted in the left panel of Fig. 9. Note that the robustness and opportuneness curves associated with the same threshold meet at $\hat{\alpha} = \hat{\beta} = 0$ at the same IR , namely the nominal performance $RR(\theta, \eta : \tilde{D}, \tilde{D}_{tot})^{-1}$.

Definition B.2. Opportune-facilitating threshold under uncertainties in the delays: The opportune-facilitating threshold θ_{RS} is the threshold that minimizes the opportuneness provided by uncertainties in the SNR that is needed for reaching the windfall performance R_w , given the SNR and nominal delays:

$$\theta_{OF}(R_w : \eta, \tilde{D}, \tilde{D}_{tot}) = \arg \min_{\theta} \hat{\beta}(R_w : \theta, \eta, \tilde{D}, \tilde{D}_{tot}). \quad (55)$$

Noting the similarity in the expression for $\hat{\alpha}$ (Eq. (46)) and $\hat{\beta}$ (Eq. (53)), the analysis of θ_{OF} follows the analysis of θ_{RS} (Eq. (48) - Eq. (50)), with the only exception that the extrema of $\hat{\beta}(R_w : \theta, \eta, \tilde{D}, \tilde{D}_{tot})$, specified by $g(\theta, \eta : R_w, \tilde{D}, \tilde{D}_{tot}) = 0$, are minima. Thus, θ_{OF} obeys the same condition, given by Eq. (23) as θ_{RS} . The resulting opportunity facilitating performance curve (OFPC) extends the RSFC given by Eq. (24) to the range of $R_w \geq \tilde{R}R(\theta, \eta : \tilde{D}, \tilde{D}_{tot})$, as depicted by the dashed curves in Fig. 3.

B.5 Performance under uncertainties in noise variance

Here we consider the performance under uncertainties in the SNR that arise from uncertainties in the noise variance. Following Appendix B.1, the performance under uncertainties in the SNR is expressed in terms of the inverse reward rate given by Eqs. (42) and (43). The worst performance possible under uncertainty α in the Info-gap model of Eq. (26) is denoted by $I_{MAX}(\alpha : \theta, \tilde{\eta}, D, D_{tot}) \equiv \left(\max_{\eta \in U(\alpha, \tilde{\eta})} IR \right)$. It follows from Eq. (43) that maximization of IR with respect to η depends on the sign of $D_{tot} - \theta$. We note that for the optimal threshold, Eq. (7) implies that $D_{tot} - \theta_{op} > 0$.

First, we consider thresholds that satisfy $D_{tot} - \theta > 0$, for which IR is maximized at minimal SNR, i.e., $\eta = \max\{0, (1 - \alpha)\tilde{\eta}\}$, so for $\alpha \leq 1$:

$$I_{MAX}(\alpha \leq 1 : \theta, \tilde{\eta}, D, D_{tot}) = \theta + D + (D_{tot} - \theta) \exp(-2\tilde{\eta}(1 - \alpha)\theta). \quad (56)$$

The case $\alpha \geq 1$ includes the worst condition of zero SNR. The resulting performance $I_{MAX}(\alpha \geq 1 : \theta, \tilde{\eta}, D, D_{tot}) = D + D_{tot}$, corresponding to instant responding at chance with $p(\text{err}) = 0.5$ (cf. Eq. (2)), is referred to as *chance performance* and denoted by $I_{chance} \equiv D + D_{tot}$. This level of performance can be achieved even under infinite uncertainty.

For thresholds that satisfy $D_{tot} - \theta < 0$, IR is maximized at maximal SNR, i.e., $\eta = (1 + \alpha)\tilde{\eta}$, so:

$$I_{MAX}(\alpha : \theta, \tilde{\eta}, D, D_{tot}) = \theta + D + (D_{tot} - \theta) \exp(-2\tilde{\eta}(1 + \alpha)\theta). \quad (57)$$

Noting that the last term is always negative, and that for $\alpha \geq 0$, the exponent is always less than unit ($\exp(-2\tilde{\eta}(1 + \alpha)\theta) < 1$) we conclude that $I_{MAX}(\alpha : \theta, \tilde{\eta}, D, D_{tot}) > \theta + D + (D_{tot} - \theta) = I_{chance}$. Sub-chance performance $I_r > I_{chance}$ is of no interest, and can be achieved even under infinite uncertainty. Thus, the analysis focuses only on thresholds of which $D_{tot} - \theta > 0$.

Considering again thresholds that satisfy $D_{tot} - \theta > 0$, we note that the minimal IR is obtained with maximal SNR $\eta = (1 + \alpha)\tilde{\eta}$, and according to Eq. (43) is given by:

$$I_{MIN}(\alpha : \theta, \tilde{\eta}, D, D_{tot}) = \theta + D + (D_{tot} - \theta) \exp(-2\tilde{\eta}(1 + \alpha)\theta). \quad (58)$$

B.6 Performance curves for uncertainties in noise variance

Here we derive RSPCs and OFPCs for uncertainties in the SNR due to uncertainties in the noise variance.

Definition B.3. Robustness to uncertainties in SNR due to uncertainties in the noise variance: *Consider a 2AFC strategy with threshold-to-drift θ and a required performance R_r . The robustness $\hat{\alpha}$ to uncertainties in the SNR is the greatest horizon of uncertainty α in the Info-gap model of Eq. (26) for which the required performance R_r can be guaranteed. The robustness is zero if the required performance*

cannot be assured even when there is no uncertainty in the SNR:

$$\hat{\alpha}(R_r : \theta, \tilde{\eta}, D, D_{tot}) = \max \left\{ \alpha \left| \left(\min_{\eta \in U(\alpha, \tilde{\eta})} RR \right) \geq R_r \vee \alpha = 0 \right. \right\}. \quad (59)$$

Performance under uncertainties in the SNR was derived in Appendix B.5 in terms of the inverse reward rate. It was concluded there that only thresholds satisfying $\theta \leq D_{tot}$ can achieve better than chance performance and thus only this range of thresholds is of interest.

Theorem B.3. Robustness curves for 2AFC under uncertainties in the SNR arising from uncertainties in the noise variance: *Consider 2AFC tasks with thresholds $\theta \leq D_{tot}$. Given the delays and nominal SNR, the robustness to uncertainties in the SNR that arise from uncertainties in the noise variance depends on the required performance R_r as:*

$$\hat{\alpha}(R_r : \theta, \tilde{\eta}, D, D_{tot}) = \begin{cases} 0, & \text{if } \tilde{R}(\theta) < R_r, \\ 1 + \frac{1}{2\tilde{\eta}\theta} \ln \left(\frac{R_r^{-1} - D - \theta}{D_{tot} - \theta} \right) & \text{if } (D_{tot} + D)^{-1} \leq R_r \leq \tilde{R}(\theta), \\ \infty, & \text{if } R_r < R_{chance}, \end{cases} \quad (60)$$

where $\tilde{R}(\theta) \equiv RR(\theta, \tilde{\eta} : D, D_{tot})$ (see Eq. (5)) is the nominal performance that can be achieved with the estimated SNR, and $R_{chance} \equiv (D + D_{tot})^{-1}$ is chance performance (see Appendix B.5).

Proof: The internal minimization of Eq. (59) is derived in Appendix B.5 by maximizing the inverse reward rate. It is shown there that chance performance $R_{chance} = (D + D_{tot})^{-1}$ can be achieved even under infinite uncertainty, so the robustness of sub-chance required performance $R_r < R_{chance}$ is infinite, as noted in the third line of Eq. (60). For better R_r , Eq. (56) implies that the internal inequality in Eq. (59) can be expressed as:

$$\theta + D + [D_{tot} - \theta] \exp(-2\tilde{\eta}(1 - \alpha)\theta) \leq R_r^{-1} \quad (61)$$

According to Eq. (5), a non-negative α can satisfy this inequality only when $R_r \leq \tilde{R}(\theta)$ where $\tilde{R}(\theta) \equiv RR(\theta, \tilde{\eta} : D, D_{tot})$ is the nominal performance. Hence, for $R_{chance} \leq R_r \leq \tilde{R}(\theta)$, the level of uncertainty that attains equality in Eq. (61) determines the robustness. Inverting Eq. (61) to express the robustness in terms of R_r proves the second line in Eq. (60). For $R_r > \tilde{R}(\theta)$, the inequality cannot be satisfied with a non-negative α , and the robustness is zero, as noted in the first line of Eq. (60). \square

Theorem B.3 indicates that the nominal performance $\tilde{R}(\theta)$, and in particular the optimal performance $\tilde{R}_{op} \equiv RR_{op}(\tilde{\eta} : D, D_{tot})$ given by Eq. (7), has zero robustness.

Definition B.4. Robust-satisficing threshold-to-drift under uncertainties in SNR: *Under uncertainties in SNR that arise from uncertainties in noise variance, the*

robust-satisficing threshold θ_{RS} that achieves the required performance R_r is the threshold that maximizes robustness to uncertainties in SNR:

$$\theta_{RS}(R_r : \tilde{\eta}, D, D_{tot}) = \arg \max_{\theta} \hat{\alpha}(R_r : \theta, \tilde{\eta}, D, D_{tot}). \quad (62)$$

Theorem B.4. Robust-satisficing performance curves (RSPCs) under uncertainties in SNR arising from uncertainties in noise variance: Denote by $q \equiv \frac{R_r^{-1} - D}{D_{tot}}$ the normalized required performance. The RSPCs, which describe the speed accuracy tradeoff implied by robust-satisficing a fixed required performance $(D + D_{tot})^{-1} \leq R_r \leq \tilde{R}_{op}$, are given by:

$$\frac{\langle DT \rangle}{D_{tot}} = \phi_{RS}(q)(1 - 2p(\text{err})) \quad \text{for } p(\text{err}) < \frac{q - \phi_{RS}(q)}{1 - 2\phi_{RS}(q) + q}. \quad (63)$$

where $\phi_{RS}(q)$ is the unique solution of:

$$\ln \left(\frac{q - \phi_{RS}}{1 - \phi_{RS}} \right) = \phi_{RS} \frac{q - 1}{(1 - \phi_{RS})(q - \phi_{RS})}. \quad (64)$$

Proof: Differentiating the robustness $\hat{\alpha}$ specified by the second line of Eq. (60) with respect to θ for a fixed level of required performance $(D + D_{tot})^{-1} \leq R_r \leq \tilde{R}_{op}$, and equating the derivative to zero, we find that the RS threshold satisfies:

$$\ln \left(\frac{R_r^{-1} - D - \theta_{RS}}{D_{tot} - \theta_{RS}} \right) = \theta_{RS} \frac{R_r^{-1} - D_{tot} - D}{(D_{tot} - \theta_{RS})(R_r^{-1} - D - \theta_{RS})}. \quad (65)$$

Denoting by $\phi_{RS} \equiv \frac{\theta_{RS}}{D_{tot}}$ the normalized RS threshold and by $q \equiv \frac{R_r^{-1} - D}{D_{tot}}$ the required performance, Eq. (65) reduces to Eq. (64).

Lemma B.1, at the end of this Appendix, shows that Eq. (64) has a unique solution, $\phi_{RS}(q)$, for $0 < q < 1$. Thus, using Eq. (2), the RSPCs for required normalized performance q , are given by Eq. (63).

We note that Eq. (65) (or the equivalent Eq. (64)) *does not* depend on the nominal SNR directly, but only indirectly via the constraint on the level of the required performance $R_r \leq \tilde{R}_{op}$. According to Eq. (7), the constraint $R_r \leq \tilde{R}_{op}$ implies that $q \geq \phi_{RS} + (1 - \phi_{RS}) \exp(-2\tilde{\eta}D_{tot}\phi_{RS})$. Using Eq. (2), this constraint reduces to the constraint for $p(\text{err})$ in Eq. (63), thereby completing the proof. \square

The resulting RSPCs, shown as solid curves in Fig. 10, are straight lines emanating from the OPC towards lower $p(\text{err})$'s and longer $\langle DT \rangle$: a more extreme leftward shift than that of Fig. 4, which does not match the higher $\langle DT \rangle$'s for mid $p(\text{err})$'s apparent in the data.

Required performance that are better than optimal, $R_r \geq \tilde{R}_{op}$, cannot be achieved robustly and are considered windfall performance (denoted as R_w). For windfall performance $R_w \geq \tilde{R}_{op}$, the opportunity-facilitating strategy can be used as detailed in Appendix B.4 for the case of uncertainties in the delays. The Opportuneness and the opportunity-facilitating (OF) threshold under uncertainties in the SNR can




Figure 10: RSPCs (solid) and OFPCs (dashed) for three levels of normalized required performance $q = (R_r^{-1} - D)/D_{tot}$ and normalized windfall performance $q = (R_w^{-1} - D)/D_{tot}$, respectively, under uncertainties in SNR arising from uncertainties in the noise variance. RSPCs and OFPCs diverge from the OPC (dot-dashed) which is obtained when $R_w = R_r = \tilde{R}_{op}(\tilde{\eta} : D, D_{tot})$.

be defined in a similar manner to the corresponding definitions under uncertainties in the delays (definitions B.1 and B.2). The opportuneness under uncertainties in the SNR depends on the best performance given by Eq. (58), instead of the worst performance given by Eq. (56). Accordingly, it can be shown that the OF threshold satisfies the same equations (65) and (64) as the RS threshold, for $R_w \geq \tilde{R}_{op}$ instead of $(D + D_{tot})^{-1} \leq R_r \leq \tilde{R}_{op}$. Thus, the OFPC (opportunity-facilitating PC) extends the RSPC for the range of large $p(\text{err})$, as demonstrated by the dashed lines in Fig. 10, with normalized windfall performance given by $q = (R_w^{-1} - D)/D_{tot}$.

Lemma B.1. For $0 < q < 1$ Eq. (64), rewritten as

$$(q - x) \ln \left(\frac{q - x}{1 - x} \right) = \frac{x(q - 1)}{(1 - x)}, \quad (66)$$

has a unique solution $x = x^*$ in the range $0 < x < q$.

Proof: The left hand side of (66) can be written as

$$f(y) = y \ln \left(\frac{y}{y + (1 - q)} \right), \quad (67)$$

where $y = q - x$. Since $(1 - q) > 0$, $f(y)$ is a decreasing function of y . The left hand side of (66) therefore increases with x from $q \ln q < 0$ at $x = 0$ to zero at $x = q$. On the other hand, the right hand side decreases monotonically from zero at $x = 0$ to $-q$ at $x = q$ (its derivative $(q - 1)/(1 - x)^2$ is negative). Both functions are continuous for $0 < x < q$, so they cross exactly once in that range. \square

B.7 Performance under uncertainties in drift rate

Here we consider uncertainties in the SNR that arise from uncertainties in the drift rate, when the noise variance is assumed known.

Proof of Theorem 3.4: Substituting the threshold-to-noise ratio $\vartheta = x_{th}/\sigma$ for the threshold-to-drift ratio $\theta = x_{th}/A$ in Eq. (43) (using $\theta = \vartheta/\sqrt{(\eta)}$), the inverse reward rate for a fixed SNR can be expressed as a function of η and ϑ :

$$IR(\vartheta, \eta : D, D_{tot}) = \frac{\vartheta}{\sqrt{\eta}}(1 - \exp(-2\sqrt{\eta}\vartheta)) + D + D_{tot} \exp(-2\sqrt{\eta}\vartheta). \quad (68)$$

For a fixed ϑ , the inverse reward rate is a decreasing function of η , and hence for presumed uncertainty $\alpha_p \leq 1$ its maximum is obtained for $\eta_w = (1 - \alpha_p)\tilde{\eta}$ and is given by:

$$I_{MAX}(\alpha \leq 1 : \vartheta, \tilde{\eta}, D, D_{tot}) = \frac{\vartheta}{\sqrt{\eta_w}} + D + (D_{tot} - \frac{\vartheta}{\sqrt{\eta_w}}) \exp(-2\sqrt{\eta_w}\vartheta). \quad (69)$$

As in Appendix B.5, the case $\alpha \geq 1$ includes the worst case of zero SNR for which the worst performance is $I_{MAX}(\alpha \geq 1 : \vartheta, \tilde{\eta}, D, D_{tot}) = D + D_{tot}$, corresponding to chance performance.

The maximin threshold, defined in 3.5 is obtained by differentiating Eq. (69) with respect to ϑ , and equating to zero. \square

B.8 Performance curves under uncertainties in drift rate

Definition B.5. Robustness to uncertainties in SNR due to uncertainties in drift rate: Consider a 2AFC strategy with threshold-to-noise ϑ and a required performance R_r . The robustness $\hat{\alpha}$ to uncertainties in the SNR is the greatest horizon of uncertainty α in the Info-gap model of Eq. (26) for which the required performance R_r can be guaranteed. The robustness is zero if the required performance cannot be assured even when there is no uncertainty in the SNR:

$$\hat{\alpha}(R_r : \vartheta, \tilde{\eta}, D, D_{tot}) = \max \left\{ \alpha \left| \left(\min_{\eta \in U(\alpha, \tilde{\eta})} RR \right) \geq R_r \vee \alpha = 0 \right. \right\}. \quad (70)$$

Definition B.6. Robust satisficing threshold-to-noise under uncertainties in SNR: Under uncertainties in the SNR that arise from uncertainties in the drift rate, the robust satisficing threshold-to-noise ϑ_{RS} for achieving the required performance R_r is the threshold that maximizes the robustness to uncertainties in the SNR:

$$\vartheta_{RS}(R_r : \tilde{\eta}, D, D_{tot}) = \arg \max_{\vartheta} \hat{\alpha}(R_r : \vartheta, \tilde{\eta}, D, D_{tot}). \quad (71)$$

Theorem B.5. Performance curves under uncertainties in SNR arising from uncertainties in the drift rate: Consider a 2AFC task with uncertainties in the SNR arising from uncertainties in the drift rate, and a better than chance required performance

$R_r > (D + D_{tot})^{-1}$. The performance curves, which specify the speed accuracy trade-off resulting from robust satisficing (for worse than optimal required performance) or opportunity-facilitating (for better than optimal desired performance) strategies are given by:

$$\frac{\langle DT \rangle}{D_{tot}} = \frac{\psi_0(1 - 2p(\text{err}))}{2 \log\left(\frac{1-p(\text{err})}{p(\text{err})}\right)}. \quad (72)$$

where ψ_0 is the solution of (with $q \equiv \frac{R_r^{-1} - D}{D_{tot}} < 1$ and $u \equiv (1 - q)/q$):

$$\exp((2 + u)\psi_0 + \sqrt{u^2\psi_0^2 + 2u\psi_0}) = \frac{1 + u\psi_0 + \sqrt{u^2\psi_0^2 + 2u\psi_0}}{q} \quad (73)$$

Proof: The internal minimization of Eq. (70) is derived in Appendix B.7 by maximizing the inverse reward rate. It is shown that for $\alpha \geq 1$, the minimal RR is the chance performance $(D + D_{tot})^{-1}$, independent of the uncertainty α . (Chance performance was defined in Appendix B.5 and as explained there can be achieved by instant choice with $p(\text{err}) = 0.5$). So if the required performance is worse than that the robustness is infinite. For better yet suboptimal R_r , the robustness satisfies the internal equality in Eq. (69):

$$\frac{\vartheta}{\sqrt{\tilde{\eta}(1 - \hat{\alpha})}} + D + [D_{tot} - \frac{\vartheta}{\sqrt{\tilde{\eta}(1 - \hat{\alpha})}}] \exp(-2\sqrt{\tilde{\eta}(1 - \hat{\alpha})}\vartheta) = R_r^{-1} \quad (74)$$

Differentiating both sides of Eq. (74) with respect to ϑ for a fixed level of required performance R_r , and equating to zero, we find that the RS threshold-to-noise, defined by the condition $\frac{d\hat{\alpha}}{d\vartheta} = 0$, satisfies:

$$\exp(2\sqrt{\tilde{\eta}(1 - \hat{\alpha})}\vartheta_{RS}) = 2D_{tot}\tilde{\eta}(1 - \hat{\alpha}) - 2\vartheta_{RS}\sqrt{\tilde{\eta}(1 - \hat{\alpha})} + 1 \quad (75)$$

Normalizing the required performance $q \equiv \frac{R_r^{-1} - D}{D_{tot}}$; we note that for better than chance performance $R_r > (D + D_{tot})^{-1}$, the normalized performance is less than 1 ($q < 1$). Eq. (75) and Eq. (74) can be combined to result in the following quadratic equation for $x \equiv \sqrt{\tilde{\eta}(1 - \hat{\alpha})}$:

$$2qD_{tot}x^2 - 2\vartheta_{RS}(q + 1)x - (1 - q - 2\vartheta_{RS}^2/D_{tot}) = 0 \quad (76)$$

whose solutions are

$$x_{1,2} = \frac{\vartheta_{RS}}{D_{tot}} \frac{q + 1}{2q} \pm \sqrt{\frac{\vartheta_{RS}^2}{D_{tot}^2} \left(\frac{1 - q}{2q}\right)^2 + \frac{1 - q}{2qD_{tot}}} \quad (77)$$

Substituting these solutions into Eq. (75), results in the following condition for the normalized threshold-to-noise $\psi \equiv \vartheta_{RS}^2/D_{tot}$:

$$\exp((2 + u)\psi \pm \sqrt{u^2\psi^2 + 2u\psi}) = \frac{1 + u\psi \pm \sqrt{u^2\psi^2 + 2u\psi}}{q} \quad (78)$$

It can be shown that with the minus sign on both sides of Eq. (78), the only solution occurs at $\psi = (1-q)/2$, leading to $x = 0$ and hence $\hat{\alpha} = 1$, which corresponds to chance performance. For better than chance performance, $\hat{\alpha} < 1$, the only relevant solution is the one corresponding to the plus sign on both sides of Eq. (78). It can be shown that such a solution exists. Denoting this solution by ψ_0 , we have: $\vartheta_{RS}^2/D_{tot} = \psi_0$. Substituting $\vartheta_{RS} = \theta_{RS}/\sqrt{\eta}$, and using Eq. (3) results in the specified performance curves of Eq. (72). When the desired performance is worse than optimal the resulting performance curve defines the RSPC, otherwise a similar analysis can show that it defines the OFPC (opportunity-facilitating performance curve) \square

The resulting performance curves, shown in Fig. 11, deviate from the experimental data in two important features: they lie below the OPCs for intermediate error rates, and reach a plateau for large error-rates close to chance.




Figure 11: Performance curves describing robust satisfying (RSPCs) and opportunity-facilitating (OFPCs) under uncertainties in drift rate given three normalized required performance levels $q = (R^{-1} - D)/D_{tot}$. The OPC is depicted by the dot-dashed curve.

B.9 Performance under uncertain and variable drift rate

Here we prove Theorem 3.5, which states that for a fixed threshold-to-noise ratio, the worst performance with a series of possibly variable SNRs, within the Info-gap model of Eq. (34), is the same as that for a fixed SNR within the Info-gap model of Eq. (26), under the same level of presumed uncertainty α_p .

Proof of Theorem 3.5: Consider a 2AFC task where the SNR at the i^{th} trial is the square-root of the i^{th} element in the infinite series $\{\eta_i\}_{i=1}^{\infty}$ obeying the Info-gap model of Eq (34). The worst SNR within the Info-gap model of Eq (34) with uncertainty level α , is given by the square-root of $\eta_w = \max(0, (1 - \alpha)\tilde{\eta})$.

Let $p(\eta : \vartheta)$ be the probability of error in a 2AFC with a fixed SNR $\sqrt{\eta}$ using the threshold-to-noise ratio ϑ . Noting that $\eta\theta = \sqrt{\eta}\vartheta$, it follows from Eq (2), that:

$$p(\eta : \vartheta) = \frac{1}{1 + \exp(2\sqrt{\eta}\vartheta)}. \quad (79)$$

When the SNR varies from trial-to-trial according to a specific series $\{\eta_i\}_{i=1}^{\infty}$, the probability of error in a 2AFC experiment of fixed duration is a random variable P_{err} , which depends on the number of trials in the experiment. Let $P(n)$ denote the probability of having n trials in the fixed duration of the experiment, the mean P_{err} (see Section A.4.3 of [20], and [14]) is:

$$\langle P_{\text{err}}(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle = \sum_{n=0}^{\infty} \frac{P(n)}{n} \left(\sum_{i=1}^n p(\eta_i : \vartheta) \right). \quad (80)$$

Since $p(\eta : \vartheta)$ given in Eq. (79) is a decreasing function of η , the mean P_{err} can only increase when replacing each η_i by η_w :

$$\langle P_{\text{err}}(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle \leq \sum_{n=0}^{\infty} \frac{P(n)}{n} \left(\sum_{i=1}^n p(\eta_w : \vartheta) \right) = p(\eta_w : \vartheta), \quad (81)$$

which is the probability of error with a fixed SNR $\sqrt{\eta_w}$.

Similarly, let $\langle DT(\eta : \vartheta) \rangle$ be the mean response time in a 2AFC with fixed the SNR $\sqrt{\eta}$ given the threshold-to-noise ratio ϑ . Noting that $\theta = \vartheta/\sqrt{\eta}$, it follows from Eq (2), that:

$$\langle DT(\eta : \vartheta) \rangle = \frac{\vartheta \exp(2\sqrt{\eta}\vartheta) - 1}{\sqrt{\eta} \exp(2\sqrt{\eta}\vartheta) + 1}, \quad (82)$$

which is a decreasing function of η (as proved in Lemma B.2 at the end of this Appendix).

The mean response time with the series of SNRs $\{\sqrt{\eta_i}\}_{i=1}^{\infty}$ is ([14]):

$$\langle DT(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle = \sum_{n=0}^{\infty} \frac{P(n)}{n} \left(\sum_{i=1}^n \langle DT(\eta_i : \vartheta) \rangle \right). \quad (83)$$

Since $\langle DT(\eta_i : \vartheta) \rangle$ given by Eq. (82) is a decreasing function of η , the mean reaction time with a series of SNRs $\langle DT(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle$ can only increase when replacing each η_i with $\sqrt{\eta_w}$:

$$\langle DT(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle \leq \sum_{n=0}^{\infty} \frac{P(n)}{n} \left(\sum_{i=1}^n \langle DT(\eta_w : \vartheta) \rangle \right) = \langle DT(\eta_w : \vartheta) \rangle, \quad (84)$$

which is the mean response time with a fixed SNR $\sqrt{\eta_w}$.

The RR , defined as the ratio between the probability of correct response divided by the average time between responses, is given by:

$$RR(\{\eta_i\}_{i=1}^{\infty} : \vartheta) = \frac{1 - \langle P_{\text{err}}(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle}{\langle DT(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle + D + (D - D_{\text{tot}}) \langle P_{\text{err}}(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle}. \quad (85)$$

Hence, the RR is a decreasing function of both $\langle P_{\text{err}}(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle$ and $\langle DT(\{\eta_i\}_{i=1}^{\infty} : \vartheta) \rangle$. According to Eqs (81) and (84), both of these factors are maximized simultaneously when the SNR is fixed at its worst level, $\sqrt{\eta_w}$. Hence, when the SNR may vary from trial-to-trial according to the series $\{\eta_i\}_{i=1}^{\infty}$ obeying the Info-gap model of Eq (34), the worst performance is obtained when $\eta_i = \eta_w$, and thus is the same as the worst performance for fixed SNR within the model of Eq. (26) for the same level of uncertainty. \square

Lemma B.2. The function $\langle DT(\eta|\vartheta) \rangle = \frac{\vartheta}{\sqrt{\eta}} \frac{\exp(2\sqrt{\eta}\vartheta)-1}{\exp(2\sqrt{\eta}\vartheta)+1}$ is a decreasing function of η

Proof: The derivative of $\langle DT(\eta|\vartheta) \rangle$ is given by:

$$\begin{aligned} & \frac{\partial \langle DT(\eta|\vartheta) \rangle}{\partial \eta} \\ &= \vartheta \left[\frac{2\vartheta \exp(2\sqrt{\eta}\vartheta) \sqrt{\eta} (\exp(2\sqrt{\eta}\vartheta)+1) - (\exp(2\sqrt{\eta}\vartheta)-1)(\exp(2\psi\vartheta)+1+2\psi\vartheta \exp(2\psi\vartheta))}{\eta (\exp(2\sqrt{\eta}\vartheta)+1)^2} \right] \\ &= \vartheta \left[\frac{-\exp(4\sqrt{\eta}\vartheta)+4\sqrt{\eta}\vartheta \exp(2\sqrt{\eta}\vartheta)+1}{\eta (\exp(2\sqrt{\eta}\vartheta)+1)^2} \right] \end{aligned} \quad (86)$$

Taylor expanding the exponent and noting that for $n > 3$, $2n < 2^n$, we may conclude that $2y \exp(y) + 1 < \exp(2y)$. Substituting $y = 2\sqrt{\eta}\vartheta$ indicates that the above derivative is negative, thereby completing the proof. \square

C The likelihood ratio test and the influence of subgroup sizes

Here we provide further details on the reliability of the statistics used in §4. We first address the assumption that the data are normally distributed. For each subgroup of subjects and each $p(\text{err})$ bin, we tested the distribution of normalized decision times for Gaussianity using the Jarque-Bera [54] and Lilliefors [55] tests, obtaining the following numbers of significantly non-Gaussian $p(\text{err})$ bins for both tests, at a significance level of 0.05. Among the top 30% and 30 – 60% groups: 3 bins in each group; among the 60 – 90% group: 2 bins; among the bottom 10% group no significantly non-Gaussian bins were found. When the 30 – 60% and 60 – 90% groups are combined into a single group as in §4, the number of significantly non-Gaussian $p(\text{err})$ bins rises to 6, probably due to the fact that some subjects in this large group employ different decision strategies, a point that we address further below. (There are 250 data points in this middle group, compared with 125 in the top 30% and 42 in the bottom 10% groups, cf. Table 1.)

When comparing two nested models, i.e., a complex model and a simple model to which the complex one can be reduced, the likelihood ratio test can be used to determine whether the data is significantly less likely under the simple rather than the complex model. This is the case when comparing the MMPC with uncertainty in D with the OPC for RR , to which the former reduces by fixing $\alpha = 0$. The likelihood

ratio test shows that the MMPC with uncertainties in D provides a significantly better description of the data for the top 30% of subjects ($p < 0.01$), and for both the other groups ($p \approx 0$). Except for the OPC for RR , all the other PCs are based on non-nested models with a single free parameter each. The corresponding likelihood ratio between the MMPC with uncertainties in D and each of these models (first four rows Table 3) provides a reasonable criterion for assessing whether the former explains the data better. A more rigorous comparison using the Bayesian approach would require additional assumptions on prior probabilities of the parameters q and α : see, e.g. [56].

As noted in §4, in fitting performance curves to data from subgroups of subjects divided according to total rewards accrued, we implicitly assume that all members of each given subgroup employ a common decision strategy. Here we probe the validity of this assumption by considering six additional ways of splitting the data. We first fitted all performance curves to the data from each subject separately and computed the average likelihood ratio for each pair of curves, as a product of likelihood ratios for the 80 individuals, obtaining the entries in the first column of Table 4. This split avoids the assumption of a common strategy, but the ratios are unreliable because individual subject data are very noisy and some bins contain as few as 3 data points, leading to overfitting. This is reflected in the fact that the entries of column 1 differ substantially from those of the other columns in Table 4.

Performance curve	Individual	20 groups	10 groups	5 groups	4 groups	2 groups
MMPC for SNR	$> 10^{13}$	$> 10^{10}$	$> 10^9$	$> 10^{10}$	$> 10^{10}$	$> 10^8$
RSPC for D	1.96	42.00	29.26	8.28	7.86	9.61
OPC for RR_m	471.54	$> 10^5$	$> 10^5$	522.09	316.03	59.16
OPC for RA	$> 10^6$	482.05	70.52	3.33	8.82	3.96
OPC for RR	$> 10^{59}$	$> 10^{33}$	$> 10^{29}$	$> 10^{26}$	$> 10^{24}$	$> 10^{20}$

Table 4: Likelihood ratios of the data given MMPC with uncertainty in D and given other performance curves (rows), for different splits of subjects into groups (columns). The higher the likelihood ratio, the less likely is it that the data could be generated by the corresponding model, in comparison to the MMPC for D . Key: MMPC: maximin performance curve; RSPC: robust satisficing performance curve; OPC: optimal performance curve; SNR: signal-to-noise ratio due to noise variance; D : delays; RR_m : modified reward-rate Eq. (13); RA : reward-accuracy rate Eq. (9); RR - reward rate Eq. (4).

This analysis was repeated five times to produce the remaining columns of Table 4, by successively dividing the subjects (sorted by total rewards accrued) into 20, 10, 5, 4, and 2 groups, with equal numbers in each group (4, 8, 16, 20, and 40 respectively). Although the precise values of the ratios depend on the number of groups, the relative ranking remains almost completely consistent. All ratios in Table 4 exceed 1, implying that the data are most likely under MMPC with uncertainty in D in all cases. The next best fits are provided by RSPC with uncertainty in D and OPC for RA , their order depending on the split. Somewhat further behind comes OPC for RR_m , trailed by MMPC with uncertainty in SNR, and finally the

parameter-free OPC for RR .

D Comparison of the pure and extended DD models for experiment 1

As explained in [22] only the first of the two experiments described in this paper yielded sufficient data for fits to the extended DD model to be feasible. Comparison of extended DD parameters obtained by such data fits, averaged over the 20 subjects, to fits of the same data to the pure DD process reveals the following. (We identify the source of the parameters by the parenthetical notes ext. and pure respectively.)

Mean values of the threshold-to-drift ratios θ (ext.) are approximately half θ (pure), mean values of SNR η (ext.) are approximately triple η (pure), while mean values of the nondecision time T_0 (ext.) are very close to T_0 (pure). In all cases the extended and pure parameter values are strongly correlated across subjects (θ : $r = 0.61, p = 0.004$, η : $r = 0.92, p < 10^{-5}$ and T_0 : $r = 0.67, p = 0.001$ [22, Fig. 4]), and the resulting optimal threshold-to-drift ratios are also strongly correlated ($r = 0.71, p < 10^{-5}$). This leads to the fact that overall correlations between subjects' thresholds and the optimal thresholds obtained from the pure and extended DD fits do not substantially differ ($r_1 = 0.44$ ($p < 10^{-5}$) for pure, and $r_1 = 0.62$ ($p < 10^{-5}$) for extended.) Moreover, comparisons of the mean thresholds (averaged over these 20 subjects) with the optimal thresholds computed analytically for the pure, and numerically for the extended DD processes, for all four delay conditions, reveal very similar patterns. See [22, Fig. 5].

We believe that this is because the substantial variances in SNR and threshold-to-drift ratio for the extended DD model are compensated by higher noise variance σ^2 in the pure DD fits (i.e., all the sources of variance are lumped in a lower SNR η (pure)). To further gauge the effects of drift and initial condition variance we compared mean decision times and error rates for the extended DD fits with the corresponding quantities evaluated for a pure DD process with η and θ equal to the mean values of the extended DD parameters. We found that $\langle DT \rangle$ s for the two processes were very similar ($r = 0.94, p < 10^{-5}$), but that $p(\text{err})$ s were significantly lower for the pure DD process (paired t-test $p < 10^{-5}$), presumably due to the unrealistically high SNRs that result when the other sources of variance are removed from the extended DD model.

Acknowledgments This work was supported by PHS grants MH58480 and MH62196 (Cognitive and Neural Mechanisms of Conflict and Control, Silvio M. Conte Center). MZ was supported by the Abramsom Center for the Future of Health and by the fund for the promotion of research at the Technion, and RB was supported by EPSRC grant EP/C514416/1. The authors thank Yakov Ben-Haim, Jonathan Cohen, Pat Simen, Angela Yu and the members of the Conte Center modeling group for numerous helpful suggestions.

References

- [1] M. Stone. Models for choice-reaction time. *Psychometrika*, 25:251–260, 1960.
- [2] D.R.J. Laming. *Information Theory of Choice-Reaction Times*. Academic Press, New York, 1968.
- [3] R. Ratcliff. A theory of memory retrieval. *Psych. Rev.*, 85:59–108, 1978.
- [4] R. Ratcliff, T. Van Zandt, and G. McKoon. Connectionist and diffusion models of reaction time. *Psych. Rev.*, 106 (2):261–300, 1999.
- [5] P.L. Smith and R. Ratcliff. Psychology and neurobiology of simple decisions. *Trends in Neurosci.*, 27 (3):161–168, 2004.
- [6] J.D. Schall. Neural basis of deciding, choosing and acting. *Nature Reviews in Neuroscience*, 2:33–42, 2001.
- [7] J.I. Gold and M.N. Shadlen. Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Science*, 5 (1):10–16, 2001.
- [8] J. Roitman and M. Shadlen. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.*, 22 (21):9475–9489, 2002.
- [9] R. Ratcliff, A. Cherian, and M.A. Segraves. A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two choice decisions. *J. Neurophysiol.*, 90:1392–1407, 2003.
- [10] M.E. Mazurek, J.D. Roitman, J. Ditterich, and M.N. Shadlen. A role for neural integrators in perceptual decision making. *Cerebral Cortex*, 13(11):891–898, 2003.
- [11] R. Ratcliff, Y.T. Hasegawa, R.P. Hasegawa, P.L. Smith, and M.A. Segraves. Dual-diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J. Neurophysiol.*, 97:1756–1774, 2006.
- [12] A. Wald. *Sequential Analysis*. Wiley, New York, 1947.
- [13] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *Ann. Math. Statist.*, 19:326–339, 1948.
- [14] J.I. Gold and M.N. Shadlen. Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36:299–308, 2002.
- [15] W. Edwards. Models for statistics, choice reaction times, and human information processing. *J. Math. Psych.*, 2:312–329, 1965.
- [16] A. Rapoport and G.B. Burkheimer. Models for deferred decision making. *J. Math. Psych.*, 8:508–538, 1971.

- [17] J.R. Busemeyer and A. Rapoport. Psychological models of deferred decision making. *J. Math. Psych.*, 32:91–134, 1988.
- [18] W.T. Maddox and C.J. Bohil. Base-rate and payoff effects in multidimensional perceptual categorization. *J. Exp. Psych.*, 24(6):1459–1482, 1998.
- [19] C.J. Bohil and W.T. Maddox. On the generality of optimal versus objective classifier feedback effects on decision criterion learning in perceptual categorization. *Memory & Cognition*, 31(2):181–198, 2003.
- [20] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J.D. Cohen. The physics of optimal decision making: A formal analysis of models of performance in two alternative forced choice tasks. *Psych. Rev.*, 113(4):700–765, 2006.
- [21] P. Holmes, E. Shea-Brown, J. Moehlis, R. Bogacz, J. Gao, G. Aston-Jones, E. Clayton, J. Rajkowski, and J.D. Cohen. Optimal decisions: From neural spikes, through stochastic differential equations, to behavior. *IEICE Transactions on Fundamentals on Electronics, Communications and Computer Science*, E88A(10):2496–2503, 2005.
- [22] R. Bogacz, P. Hu, J.D. Cohen, and P. Holmes. Do humans select the speed-accuracy tradeoff maximizing reward rate? *Cognitive, Affective and Behavioral Neuroscience*, ??:??, 2008. submitted.
- [23] Y. Ben-Haim. *Information Gap Decision Theory: Decisions under Severe Uncertainty*. Academic Press, New York, 2006. 2nd Edition.
- [24] H.A. Simon. *Models of bounded rationality*. MIT Press, Cambridge, MA, 1982.
- [25] A.S. Reber. *Practical Nonparametric Statistics*The Penguin Dictionary of Psychology. Penguin Books, London, U.K., 1995.
- [26] C.W. Gardiner. *Handbook of Stochastic Methods, Second Edition*. Springer, New York, 1985.
- [27] J.R. Busemeyer and J.T. Townsend. Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100:432–459, 1993.
- [28] R. Ratcliff and F. Tuerlinckx. Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonom. Bull. Rev.*, 9:438–481, 2002.
- [29] J. Ditterich. Evidence for time-variant decision making. *European J. of Neurosci.*, 24:3628–3641, 2006.
- [30] J. Ditterich. Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks*, 19:981–1012, 2006.

- [31] Y. Liu, P. Holmes, and J.D. Cohen. A neural network model of the Eriksen task: Reduction, analysis, and data fitting. *Neural Computation*, 20 (2):345–373, 2008.
- [32] P. Eckhoff, P. Holmes, C. Law, P.M. Connolly, and J.I. Gold. On diffusion processes with variable drift rates as models for decision making during learning. *New J. of Physics*, 10:doi:10.1088/1367-2630/10/1/015006, 2008.
- [33] M. Usher and J.L. McClelland. On the time course of perceptual choice: The leaky competing accumulator model. *Psych. Rev.*, 108:550–592, 2001.
- [34] J.D. Cohen, K. Dunbar, and J.L. McClelland. On the control of automatic processes: A parallel distributed processing model of the Stroop effect. *Psychological Review*, 97(3):332–361, 1990.
- [35] X.-J. Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36:955–968, 2002.
- [36] K.F. Wong and X.-J. Wang. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.*, 26 (4):1314–1328, 2006.
- [37] P. Simen, J.D. Cohen, and P. Holmes. Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks*, 19:1013–1026, 2006.
- [38] K.H. Britten, M.N. Shadlen, W.T. Newsome, and J.A. Movshon. Responses of neurons in macaque MT to stochastic motion signals. *Visual Neurosci.*, 10:1157–1169, 1993.
- [39] I.J. Myung and J.R. Busemeyer. Criterion learning in a deferred decision making task. *Amer. J. Psych.*, 102(1):1–16, 1989.
- [40] J.R. Busemeyer and I.J. Myung. An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *J. Exp. Psych. General*, 121(2):177–194, 1992.
- [41] H.A. Simon. Rational choice and the structure of environments. *Psych. Rev.*, 23:129–138, 1956.
- [42] G. Gigerenzer. Decision making: nonrational theories. In N.J. Smelser and P.B. Bates, editors, *International Encyclopedia of the Social and Behavioral Sciences*, volume 4, pages 3304–4409. MIT Press, Cambridge, MA, 2001.
- [43] G. Gigerenzer. The adaptive toolbox. In G. Gigerenzer and R. Selten, editors, *Bounded Rationality: the adaptive toolbox*. MIT Press, Cambridge, MA, 2002.
- [44] G. Gigerenzer and D.G. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psych. Rev.*, 103:650–669, 1996.
- [45] D. Ellsberg. Risk, ambiguity and the Savage axioms. *Quart. J. Economics*, 75:643–669, 1961.

- [46] Y. Carmel and Y. Ben-Haim. Info-gap robust-satisfying model of foraging behavior: Do foragers optimize or satisfice? *American Naturalist*, 166(5):633–641, 2005.
- [47] J. Gibbon. Scalar expectancy theory and Weber’s law in animal timing. *Psych. Rev.*, 84(3):279–325, 1977.
- [48] C.V. Buhusi and W.H. Meck. What makes us tick? functional and neural mechanisms of interval timing. *Nature*, 6:755–765, 2005.
- [49] H. Akaike. Likelihood of a model and information criteria. *J. Econometrics*, 16:3–14, 1981.
- [50] H. Lujendijk. Practical experiment on noise perception in noisy images. *Proc. Soc. of Optical Engineering*, 2166:2–8, 1994.
- [51] P Frazier and A.J. Yu. Sequential hypothesis testing under stochastic deadlines. In *Advances in Neural Information Processing Systems*. Neural Information Processing Systems Foundation, 2007. Downloadable from <http://books.nips.cc/nips20.html>.
- [52] D.R. Cox. *Renewal Theory*. Methuen, London, 1962. Monographs on applied probability and statistics.
- [53] M. Abramowitz and I.A. Stegun, eds. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, New York, 1972.
- [54] G.G. Judge, R.C. Hill, W.E. Griffiths, H. Lutkepohl, and T.-C. Lee. *Introduction to the Theory and Practice of Econometrics*. Wiley, New York, 1988.
- [55] W.J. Conover. *Practical Nonparametric Statistics*. Wiley, New York, 1980.
- [56] D.J.C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, Cambridge, U.K., 2003.