RESEARCH ARTICLE

Probabilistic Decision Making with Spikes: From ISI Distributions to Behaviour via Information Gain

Javier A. Caballero[1,2], Nathan F. Lepora[3,4], Kevin N. Gurney[1] *

1 Dept of Psychology, University of Sheffield, Sheffield, UK, 2 Faculty of Life Sciences, University of Manchester, Manchester, UK, 3 Dept of Engineering Mathematics, University of Bristol, Bristol, UK, 4 Bristol Robotics Laboratory, University of Bristol and University of the West of England, Bristol, UK

* k.gurney@shef.ac.uk

Abstract

Computational theories of decision making in the brain usually assume that sensory 'evidence' is accumulated supporting a number of hypotheses, and that the first accumulator to reach threshold triggers a decision in favour of its associated hypothesis. However, the evidence is often assumed to occur as a continuous process whose origins are somewhat abstract, with no direct link to the neural signals - action potentials or 'spikes' - that must ultimately form the substrate for decision making in the brain. Here we introduce a new variant of the well-known multi-hypothesis sequential probability ratio test (MSPRT) for decision making whose evidence observations consist of the basic unit of neural signalling - the inter-spike interval (ISI) - and which is based on a new form of the likelihood function. We dub this mechanism s-MSPRT and show its precise form for a range of realistic ISI distributions with positive support. In this way we show that, at the level of spikes, the refractory period may actually facilitate shorter decision times, and that the mechanism is robust against poor choice of the hypothesized data distribution. We show that s-MSPRT performance is related to the Kullback-Leibler divergence (KLD) or information gain between ISI distributions, through which we are able to link neural signalling to psychophysical observation at the behavioural level. Thus, we find the mean information needed for a decision is constant, thereby offering an account of Hick's law (relating decision time to the number of choices). Further, the mean decision time of s-MSPRT shows a power law dependence on the KLD offering an account of Piéron's law (relating reaction time to stimulus intensity). These results show the foundations for a research programme in which spike train analysis can be made the basis for predictions about behavior in multi-alternative choice tasks.

Introduction

The decisions we make every day rely on processing continually refreshed streams of uncertain information. This information guides our choices, until some assumed termination criterion is

reached, upon which a decision is made [1–6]. Previous influential frameworks have addressed
the concerns of uncertainty and time continuity by assuming that new information or *evidence*
occurs as continuous stochastic processes [1, 7–10], often mapped at the level of the membrane
potential of individual neurons [11–14]. However, most of the ensuing mechanisms are not
naturally suited to exploit the statistical structure of the discrete and irregular sequences of ac-
tion potentials that must ultimately form the substrate for decision making in the brain. As no-
table exceptions, [15–17] assumed that evidence is sampled from Poisson processes and that
statistical inference is conducted upon them. However, their result is founded on the *assump-
tion* that, for Poisson-based spike trains, the evidence was given by counting spikes, and it is
not apparent what theoretical foundation might underly this assumption.

Further, there is, as yet, no clear way of dealing with spike trains with arbitrary inter-spike-
interval (ISI) statistics, although there has been some interest in log-ISIs as means for sampling
for decision formation [18]. Almost all previous models of decision making have worked at a
more abstract behavioural level, where the interpretation of evidence is less constrained.
Hence, they have either assumed the evidence is Gaussian distributed (*e.g.* [1, 7, 8, 10]) or they
have used a general formalism with no distribution specified (*e.g.* [19]). Both views have issues:
the negative tail of the Gaussian makes it conceptually inadequate for describing ISIs (or their
inverse, the instantaneous firing rate) and better fits are to be had with other skewed probabili-
ty density functions (pdfs); while the general view does not provide any connection with neural
mechanisms, in particular spikes. Typical neural recordings show spike trains which cannot be
described by Poisson processes (*e.g.* see [11, 20–27]) and these should be amenable to the same
theoretical analysis as their Poisson counterparts. Fortunately, there has been considerable at-
tention paid to identifying the particular pdf that adequately describes variables like these [20,
22, 23, 26, 27]: it is known that the typical distribution of ISIs is positively skewed and has a
non-zero mode (maximum). Any neurally grounded account of decision making must accom-
modate these data.

To address these issues, we introduce here, a general approach to decision making with
spikes based only on the assumption that evidence for decisions is conveyed in the distribution
of ISIs. Thus, our premise is that the arrival of spikes provides the primary basis of information
transfer and that the ISI provides the basic 'unit of evidence'. The ISI is therefore considered to
be an 'atomic' time interval that is not linked to some more fundamental, continuous-time pro-
cess (as is often the case). Our proposal gravitates around a novel variant of the likelihood func-
tion which can be applied to any Bayesian or frequentist inference algorithm that uses
sequential sampling (for excellent reviews see [9, 28]). With this likelihood we produce a new
principled and general formulation of a statistical inference procedure known as the multi-hy-
pothesis sequential probability ratio test [29] (MSPRT). We dub this mechanism *spiking
MSPRT* or simply s-MSPRT.

This novel and more general spike-based account allows us to address several previously un-
answered questions. This was made possible by investigating several particular cases of the s-
MSPRT, each of which assumed a pdf compatible with the distribution of ISIs in existing data
[30, 31]. First, noting that the non-zero mode in typical ISI distributions is dictated by an inter-
spike refractory period, we ask, what is the advantage of such shape? The refractory period is
often seen as drawback for information transmission, but we show it actually furnishes ISIs
with additional information for discrimination, thereby facilitating their ability to inform
decision making.

Second, given the diversity of s-MSPRT instances now possible, how robust is decision mak-
ing if the distributions tested-for (the basis of the particular s-MSPRT) are not the same as

those of the data being tested? Hitherto, this issue has not been as acute because of the assumption of Gaussian signals and test-distributions. However, as well as parametric differences between test and data distributions, we now have to consider the possibility of a fundamental mismatch between the two, reflecting a potential 'ignorance' of the algorithms regarding the actual statistical structure of the task [32]. Here, we show that this issue does not fundamentally compromise the s-MSPRT, thereby revealing a robustness we might expect of a biological mechanism.

Finally we ask, what are the implications of positing a *neural-level* mechanism for decision making like s-MSPRT, at the behavioural level? Our starting point for this analysis is the observation that the mean decision sample size (mean *decision sample* for short) of our algorithms is intimately related to the discrimination information between the distributions of the input data streams, as measured by the Kullback-Leibler divergence (KLD) [33, 34]. Based on this, we demonstrate that the mean total information needed to reach any particular decision is constant and thus conserved. We show how this postulate gives rise to an expression consistent with Hick's law [35], a well-known psychometric regularity where mean reaction times are shown to approximate a logarithmic function of the number of choices. Lastly, from the same postulate we demonstrate that the mean decision sample of any of our algorithms depends through a power law on the KLD. This bears a strong resemblance to a second regularity known as Piéron's law [36], where the mean reaction time decreases following a power law as the intensity of the stimulus increases. The KLD between the possible ISI distributions of a sensory neuron should increase for more intense stimuli, we argue, therefore, that our postulate on the conservation of mean total information is a possible explanation for Piéron's law. We conclude that, the hypothesis that the brain approximates an algorithm like s-MSPRT in its decision making, is consistent with several behavioural phenomena, and that the s-MSPRT provides an explanation for these phenomena grounded in information theory.

Earlier versions of some of these results have been reported in conference abstracts [37, 38] and JC's PhD dissertation [39].

Results

s-MSPRT: a Bayes-based decision mechanism using spikes

We establish an idealised decision making mechanism grounded in the Bayesian approach, which works by sampling a set of spike trains encoding information or *evidence* about a stimulus in a set of parallel data streams or *channels*. The mechanism is a variant of the multi-hypothesis sequential probability ratio test (MSPRT) [29] and must decide which of a set of hypotheses about these data streams is the most probable.

Sequential sampling algorithms usually assume that all data streams are sampled synchronously, and evidence is thereby accumulated for all hypotheses simultaneously. However, our point of departure is to suppose that, in a neural context, evidence is supplied on each channel of a decision mechanism by the arrival of spikes on that channel. Thus, there is a notion of sampling in spike trains grounded in their very construction—that the arrival of a spike within a channel supplies new evidence therein, and conversely, that no new data for the channel is supplied between its spike arrivals. The implication of this is that the notion of synchronous or uniform sampling is lost, because, in general, spike arrivals across channels will be asynchronous. However, it is still possible to establish a sequential sampling scheme that can be used in Bayesian (and indeed, frequentist) inference.

To proceed, consider a decision process over C data channels, each one comprised of a neural spike train. Let the process start at t^0, and let $t_i^0, t_i^1, \ldots, t_i^j, \ldots$ be the sequence of spike arrival

times on channel i, where $t_i^0 \geq t^0$, and $t_i^j > t_i^{j-1}$ for all j. We assume the basic unit of data being supplied to the neural decision mechanism is the inter-spike-interval (ISI) defined by $x_i(j) = t_i^j - t_i^{j-1}$. The ISIs also serve to define the (non-uniform) data 'sampling times' or 'sampling intervals' on a *per-channel* basis. There is, therefore, no way of exactly assigning a number of observations consistently across all channels up to a specified time T. To define the channel-specific sample size, let $s_i(T) = \text{argmax}_j\{t_i^j \leq T\}$. The first and last spikes in the interval $[t^0, T]$ for channel i occurred at t_i^0, and $t_i^{s_i(T)}$ respectively; note that time t^0 does not necessarily coincide with any spike arrival. Then, the decision process is effective on this channel for an interval $T_i = t_i^{s_i(T)} - t_i^0$, and there are $s_i(T)$ ISIs on i during this time. Now let $\mathbf{x}_i(T)$ be the set of observations (equivalently ISIs) for channel i during the decision process, where $\mathbf{x}_i(T) = (x_i(1), x_i(2),\ldots x_i(s_i(T)))$. Finally, let $X(T) = \{\mathbf{x}_k(T), k = 1,\ldots, C\}$. Thus, $X(T)$ is the entire data set available to the decision mechanism up to time T.

We now suppose there are set of N hypotheses H_i, with $i \in \{1,\ldots, N\}$, about the data which we wish to test. At this stage, we retain the most general formalism in which the number of hypotheses, N, is not necessarily the same as the number of channels, C [40]. Further the hypotheses may concern perceptual interpretations of the data rather than their statistical properties as such [19]. However, subsequently, we specialise to the case where the number of data channels and hypotheses is the same and make more precise the nature of the hypotheses themselves. In general, the hypothesis test requires we compute the posterior probabilities $P(H_i|X(T))$. Using Bayes rule we have

$$P(H_i|X(T)) = \frac{P(X(T)|H_i)P(H_i)}{\sum_{k=1}^{N} P(X(T)|H_k)P(H_k)} \tag{1}$$

where $P(H_i)$ are the priors for each hypothesis, and $P(X(T)|H_i)$ the likelihoods. While our general model allows for a manipulation of the priors to bias choices, for simplicity we hereafter assume that they are all equal. We also take logarithms, thereby transforming fractions to sums

$$\log P(H_i|X(T)) = \log P(X(T)|H_i) - \log \sum_{k=1}^{N} P(X(T)|H_k) \tag{2}$$

Now, following [10, 19], we write this as

$$\log P(H_i|X(T)) = \log P(X(T)|H_i) - \log \sum_{k=1}^{N} \exp\left(\log P(X(T)|H_k)\right) \tag{3}$$

(Lepora and Gurney [19] actually dealt with the negative of the log-posterior which allows the interpretation that the decision is performed by basal ganglia [10]; this is not essential to our exposition here). Then putting $P_i(T) \equiv P(H_i|X(T))$ and $y_i(T) \equiv \log P(X(T)|H_i)$

$$\log P_i(T) = y_i(T) - \log \sum_{k=1}^{N} \exp\left(y_k(T)\right) \tag{4}$$

It is apparent that a key computation here is the log-likelihood $y_i(T)$. Assuming independence of data across channels, and no dependence between ISIs within a channel

$$
\begin{aligned}
y_i(T) &= \sum_{k=1}^{C}\sum_{j=1}^{s_k(T)} \log p(x_k(j)|H_i) \\
&= \sum_{j=1}^{s_i(T)} \log p(x_i(j)|H_i) + \sum_{\substack{k=1 \\ k \neq i}}^{C} \sum_{m=1}^{s_k(T)} \log p(x_k(m)|H_i)
\end{aligned}
\tag{5}
$$

Where $p(x_k(j)|H_i)$ is some probability measure applied to $x_k(j)$ under the hypothesis H_i, e.g. a probability density (from a pdf) or probability (from a probability mass function). Thus far, we have a very general situation where an arbitrary number of data streams or channels C (e.g. spike trains from individual neurons) can contribute to any number of hypotheses N. However, we now specialise to the case when there is the same number of hypotheses as data streams ($C = N$), and each hypothesis H_i takes the following form: that the i.i.d. data $x_i(j)$ in channel i was drawn from a 'preferred' pdf f_*, with mean μ^* and standard deviation σ^*, while i.i.d. data in other channels $x_k(j)$, $k \neq i$ were drawn from a 'null' pdf f_0 with mean μ_0 and standard deviation σ_0 (in general, $\mu^* \neq \mu_0 \neq \sigma^* \neq \sigma_0$). Thus

$$
y_i(T) = \sum_{j=1}^{s_i(T)} \log f_*(x_i(j)) + \sum_{\substack{k=1 \\ k \neq i}}^{N} \sum_{m=1}^{s_k(T)} \log f_0(x_k(m))
\tag{6}
$$

The form of the log-likelihood may be simplified by expressing it in terms of probability ratios [10, 19]. To do this we rewrite Eq 6 as

$$
y_i(T) = \sum_{j=1}^{s_i(T)} \log f_*(x_i(j)) - \sum_{j=1}^{s_i(T)} \log f_0(x_i(j)) + \sum_{k=1}^{N}\sum_{m=1}^{s_k(T)} \log f_0(x_k(m))
\tag{7}
$$

The double sum on the extreme right is hypothesis independent and we denote it by $B(T)$. Then

$$
y_i(T) = \sum_{j=1}^{s_i(T)} \log \frac{f_*(x_i(j))}{f_0(x_i(j))} + B(T)
\tag{8}
$$

Note that this novel definition of the likelihood, grounded directly in the ISIs of the spike trains, is quite general and may be also used in frequentist sequential sampling methods. It is straightforward to show using Eq 4 that hypothesis independent terms like $B(T)$ do not alter the posterior, and so we only need consider the first term in Eq 8. We therefore redefine $y_i(T)$ to be this term, and also introduce some additional notation

$$
y_i(T) = \sum_{j=1}^{s_i(T)} L_i(j) , \qquad L_i(j) = \log \frac{f_*(x_i(j))}{f_0(x_i(j))}
\tag{9}
$$

which defines the per-observation ratios $L_i(j)$. These constitute the 'evidence contributions', so that $y_i(T)$ is the 'accumulated evidence'.
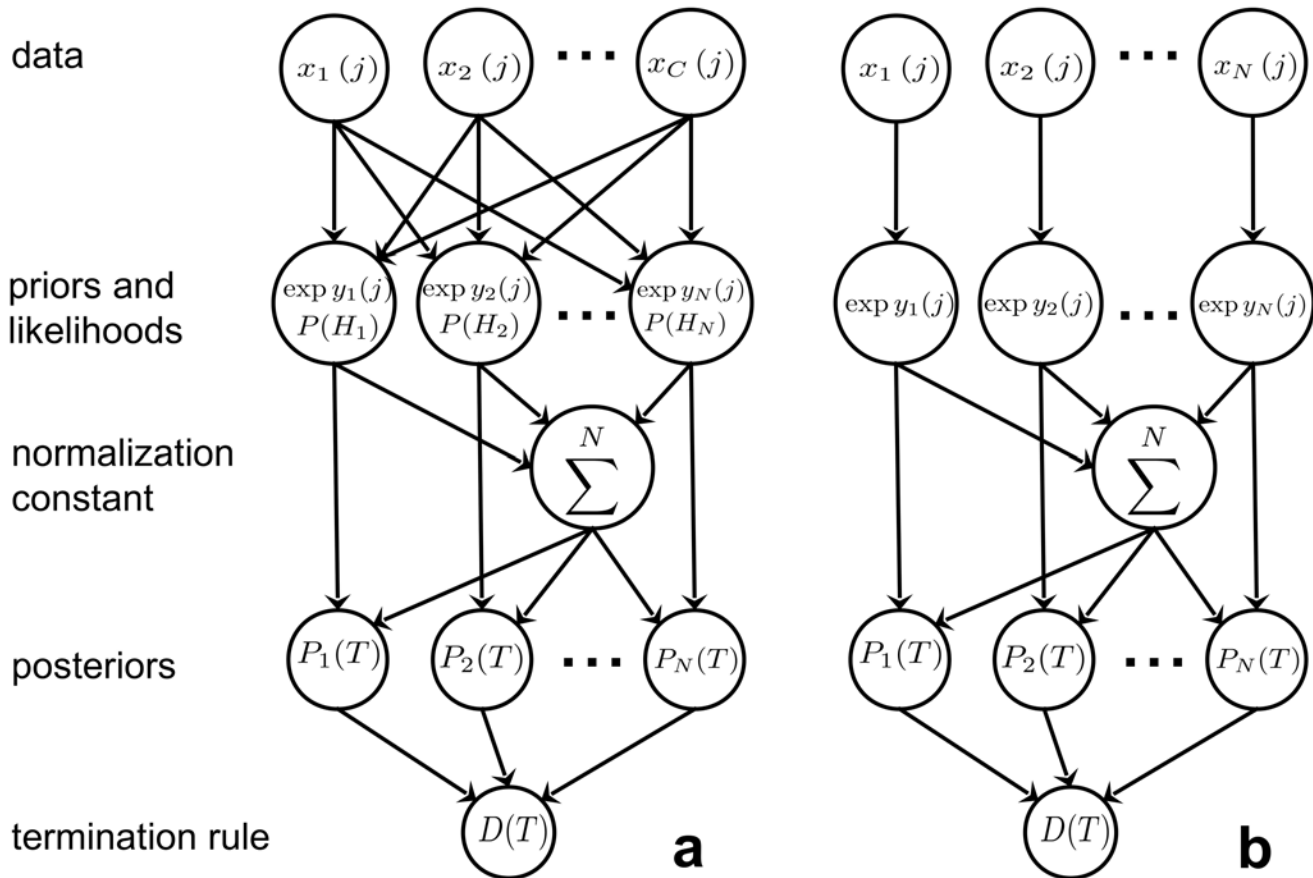
Fig 1. MSPRT in schematic form. Panel a shows the general MSPRT where all the *C* data streams contribute to all of the *N* likelihoods and thus posteriors, which are then evaluated at a termination stage. Panel b only shows the effective components after all simplifications have been applied.

Then, after substituting Eq 9 in Eq 4, the decision *D*(*T*), at time *T* is:

$$D(T) = \begin{cases} \text{Choose hypothesis } i: & \text{if } \log P_i(T) = \max_{j \in \{1,...,N\}} \log P_j(T) \geq \theta, \text{at } T = T_D \\ \text{Continue sampling}: & \text{if } \max_{j \in \{1,...,N\}} \log P_j(T) < \theta \end{cases} \quad (10)$$

Although an individual threshold can be set per hypothesis [29], for simplicity [10, 16, 17] we assume the single $\theta \in [\log(1/N),0)$; where the position of θ controls the speed-accuracy trade-off. Informally, Eq 10 states that the decision at any given time is: either picking the most salient (likely) hypothesis (*i*), if its decision variable ($\log P_i(T)$) has surpassed a threshold θ at the spike arrival time T_D, or continuing to gather data. In what follows, Eq 10 is called the *spiking MSPRT* or simply *s-MSPRT*. In Fig 1 it is shown in schematic form to illustrate the flow of information through the algorithm.

Fig 2 shows the time course of a single trial in s-MSPRT which serves to illustrate several key points. First, unlike several other instances of the use of MSPRT in neural decision making, s-MSPRT has its sampling grounded in a physically observable process—spike arrivals—which ties it directly to time. Thus, there is no arbitrary synchronous sampling time of some more
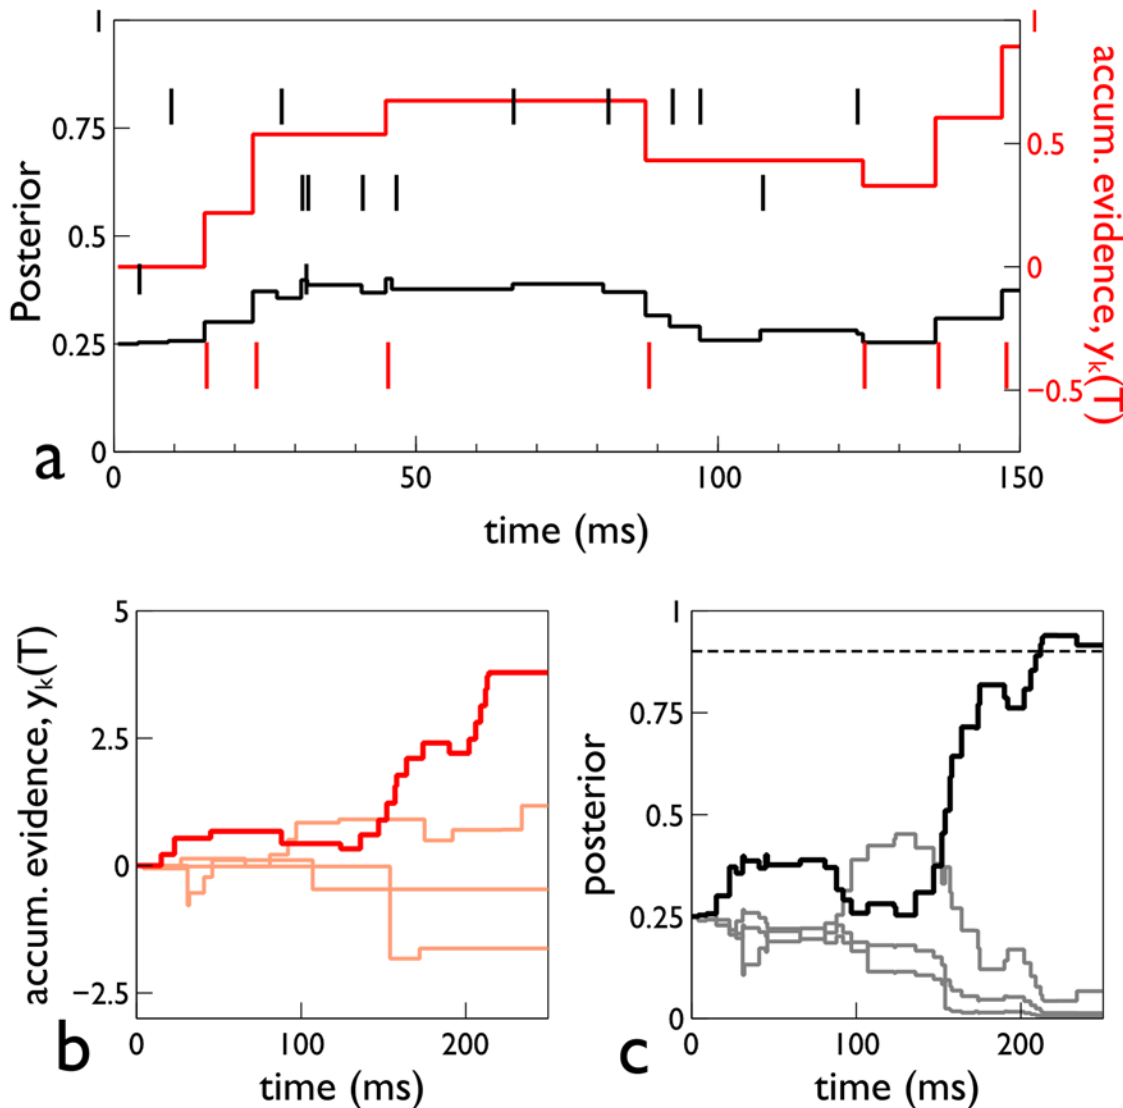
Fig 2. Time course of signals in a single trial of s-MSPRT. The trial is for an s-MSPRT using the gamma distribution, with 4 choices, under the parameterisation set Ω_{IV} (see Methods). Panel a, shows the spike rasters of the 4 spike trains as small vertical line markers, with that of the preferred channel in red. This panel also shows the accumulated evidence $y_k(T)$ as a red line graph. Panel b shows $y_k(T)$ of all four hypotheses with the preferred hypothesis in bold red. Panel c shows the posteriors with the preferred hypothesis in black and the others in gray. The threshold is shown by the horizontal dashed line, and was chosen to give a 5% error rate.

fundamental continuous-time process, and the decision time emerges naturally in terms of the input data streams. A second observation is that sampling is not uniform; for a given time T, there are different numbers of observations per channel, $s_i(T)$, which depend on the firing in that channel. This means the $y_i(T)$ are updated at different times, see Fig 2b. Third, while the contribution $y_i(T)$ is updated only at spike arrival times on channel i, the log-posterior $\log P_i(T)$ is updated at the arrival of a spike on *any* channel (see Eq 4 and black line in Fig 2a). Finally, the preceding analysis is quite general; no requirements have been made on the form of the distributions of ISIs (Gaussian or otherwise). In the next two sub-sections, we go on to consider the consequences of using different neurobiologically plausible forms for f_0, f_*.

Accumulated evidence with exponentially distributed ISIs is approximately a spike-count

Consider an s-MSPRT with hypotheses that the data is distributed exponentially $f_*(x) = \lambda_* \exp(-\lambda_* x), f_0(x) = \lambda_0 \exp(-\lambda_0 x)$. Here, λ_* and λ_0 are mean instantaneous firing rates, defined by the inverse of the respective mean ISI μ^*, μ_0. Then, substituting these forms in Eq 9.

$$
\begin{aligned}
y_i(T) &= \sum_{j=1}^{s_i(T)} \left[\log\left(\frac{\lambda_*}{\lambda_0}\right) - x_i(j)(\lambda_* - \lambda_0) \right] \\
&= s_i(T) \log\left(\frac{\lambda_*}{\lambda_0}\right) - T_i(\lambda_* - \lambda_0) \\
&= s_i(T) \log\left(\frac{\lambda_*}{\lambda_0}\right) - T(\lambda_* - \lambda_0) + \delta_i(T)
\end{aligned}
\tag{11}
$$

where $\delta_i(T) = (T - T_i)(\lambda_* - \lambda_0)$ and we have used the fact that summing consecutive ISIs just gives their overall duration (here, T_i). The term $T(\lambda_* - \lambda_0)$ is hypothesis independent and can be absorbed into $B(T)$ in Eq 8. We therefore have

$$
y_i(T) = s_i(T) \log\left(\frac{\lambda_*}{\lambda_0}\right) + \delta_i(T)
\tag{12}
$$

or

$$
y_i(T) \approx g s_i(T), \text{ with } g = \log\left(\frac{\lambda_*}{\lambda_0}\right)
\tag{13}
$$

and where the 'error' term defining the degree of approximation is $\delta_i(T)$.

The right hand side of Eq 13 is a spike count scaled by a 'gain' $g = \log(\lambda_*/\lambda_0)$. This approximation resembles the expression used by Zhang & Bogacz [17] for $y_i(T)$. The difference between the precise form for $y_i(T)$ here (Eq 12) and that of [17] stems from them formulating their likelihood function upon discrete Poisson probability mass functions (of spike counts) and us in terms of continuous exponential pdfs (assumed here for ISIs). Nevertheless, up to the approximation in Eq 13, we concur with Zhang & Bogacz that, for Poisson-based spike trains (with exponential ISI distributions), the accumulated evidence $y_i(T)$ is given by a (gain-multiplied) spike count. However, in deriving their result, Zhang & Bogacz started by assuming that $y_i(T)$ comprised the spike count, and then deduced and included a separate gain factor. In contrast, we have not assumed any a priori form for the total evidence, and have obtained the integrated form (counts and gain) in Eq 12 directly from the general expression for the sum of individual spike contributions in Eq 8.

To obtain a better understanding of the spike count approximation in Eq 13, use the gain defined there in Eq 12 to obtain

$$
y_i(T) = g(s_i(T) + \hat{s}_i(T))
$$

where $\hat{s}_i(T) = \delta_i(T)/g$ is the error in the spike count. Now consider the spikes immediately prior to, and after T. The expectation of the intervals between these spikes and T are equal, and their sum is the mean ISI μ. Therefore $\langle T - T_i \rangle = 0.5\mu$ and so, for the preferred channel

$$
\begin{aligned}
\langle \hat{s}_i(T) \rangle &= \frac{0.5}{g} \left(1 - \frac{\lambda_0}{\lambda_*} \right) \\
&= \frac{0.5}{g} (1 - e^{-g})
\end{aligned}
\tag{14}
$$

Table 1. Analytic expressions for 'evidence contributions', $L_i(j)$ for a range of distributions.

distribution $f(x)$	$L_i(j)$	gains
exponential $\lambda \exp(-\lambda x)$	$g_0^E - g_1^E x_i(j)$	$g_0^E = \log\left(\frac{\lambda_*}{\lambda_0}\right)$ $g_1^E = \lambda_* - \lambda_0$
lognormal $\frac{1}{x\Theta\sqrt{2\pi}}\exp\left[\frac{-(\ln x - \kappa)^2}{2\Theta^2}\right]$	$g_0^L + g_1^L[\log x_i(j)]^2 + g_2^L\log x_i(j)$	$g_0^L = \frac{\kappa_0^2}{2\Theta_0^2} - \frac{\kappa_*^2}{2\Theta_*^2} + \log\left(\frac{\Theta_0}{\Theta_*}\right)$ $g_1^L = \frac{1}{2\Theta_0^2} - \frac{1}{2\Theta_*^2}$ $g_2^L = \frac{m_*}{\Theta_*^2} - \frac{m_0}{\Theta_0^2}$
gamma $\frac{x^{\alpha-1}}{\Gamma(k)\beta^z}\exp\left(\frac{-x}{\beta}\right)$	$g_0^\gamma + g_1^\gamma\log x_i(j) + g_2^\gamma x_i(j)$	$g_0^\gamma = \log\left[\frac{\Gamma(\alpha_0)\beta_0^{z_0}}{\Gamma(\alpha_*)\beta_*^{z_*}}\right]$ $g_1^\gamma = \alpha_* - \alpha_0$ $g_2^\gamma = \frac{1}{\beta_0} - \frac{1}{\beta_*}$
inverse gamma $\frac{x^{-(\varphi+1)}}{\rho^\varphi\Gamma(\varphi)}\exp\left(-\frac{1}{\rho x}\right)$	$g_0^M + g_1^M\log x_i(j) + g_2^M x_i(j)^{-1}$	$g_0^M = \log\left[\frac{\rho_0^{\varphi_0}\Gamma(\varphi_0)}{\rho_*^{\varphi_*}\Gamma(\varphi_*)}\right]$ $g_1^M = \varphi_0 - \varphi_*$ $g_2^M = \frac{1}{\rho_0} - \frac{1}{\rho_*}$
inverse Gaussian $\left(\frac{v}{2\pi x^3}\right)^{1/2}\exp\left[\frac{-v(x-\mu)^2}{2\mu^2 x}\right]$	$g_0^S + g_1^S x_i(j)^{-1} + g_2^S x_i(j)$	$g_0^S = \frac{1}{2}\log\left(\frac{v_*}{v_0}\right) + \frac{v_*}{\mu_*} - \frac{v_0}{\mu_0}$ $g_1^S = \frac{v_0 - v_*}{2}$ $g_2^S = \frac{1}{2}\left(\frac{v_0}{\mu_0} - \frac{v_*}{\mu_*}\right)$

The left hand column shows the canonical functional form of the distribution in terms of its 'natural' parameters. The central column is the evidence contribution $L_i(j)$, and the right hand column contains expressions for the 'gains' in terms of such parameters.

This takes its largest values as $g \to 0$, in which case $\langle \hat{s}_i(T)\rangle \to 0.5$. There is therefore an upper bound on the expected error in the spike count to decision on the preferred channel of 0.5.

In general accumulated evidence is not given by a spike-count alone

The exponential distribution is not privileged in its relation to Eq 9. Thus, it is possible to obtain a closed form for $L_i(j)$ for any analytically defined distribution, by substituting a suitably parameterised pair of pdfs for $f_*(x)$ and $f_0(x)$ into Eq 9. The results for a range of distributions which may fit ISI data are shown in Table 1. Each distribution has a pair of parameters ζ, η (not usually the mean and standard deviation of the pdf). Taking two pairs of such parameters ζ_*, η_* and ζ_0, η_0, specifies $f_*(x)$ and $f_0(x)$. Substitution in $L_i(j)$ in Eq 9 yields the expressions in the central column of the table. Each one comprises the sum of a constant g_0, and a sum of products of a coefficient or 'gain', g_i ($i = 1,2$), together with a simple function of the variable like $\log x_i(j)$, $(\log x_i(j))^2$, $x_i(j)^{-1}$. Summation over spikes leads, in all instances, to a term like $g_i s_k(T)$, which is the analogue of the right hand side of Eq 13; that is, it expresses a 'spike count' contribution to $y_i(T)$.

However, there are other terms in $L_i(j)$ which depend on $x_i(j)$ and which can make a substantial contribution to the development of $y_i(T)$. This is demonstrated in Fig 3 for a particular set of biologically plausible pdfs and parameterisations, Ω_{IV}, described in the Methods. Fig 3a–3d, shows the pdfs specified by Ω_{IV}. In the corresponding panels e-h below, $L_i(j)$ is shown per pdf as a function of $x_i(j)$ (solid red line) as well as the functions for its contributory terms. The term g_0 is constant (solid black line), the spike count contribution and other non-linear terms are functions of $x_i(j)$ (dashed and broken gray/black lines). There is clearly a wide variation in non-constant contributions to $L_i(j)$. Most notably, the terms linear in $x_i(j)$ in the gamma and
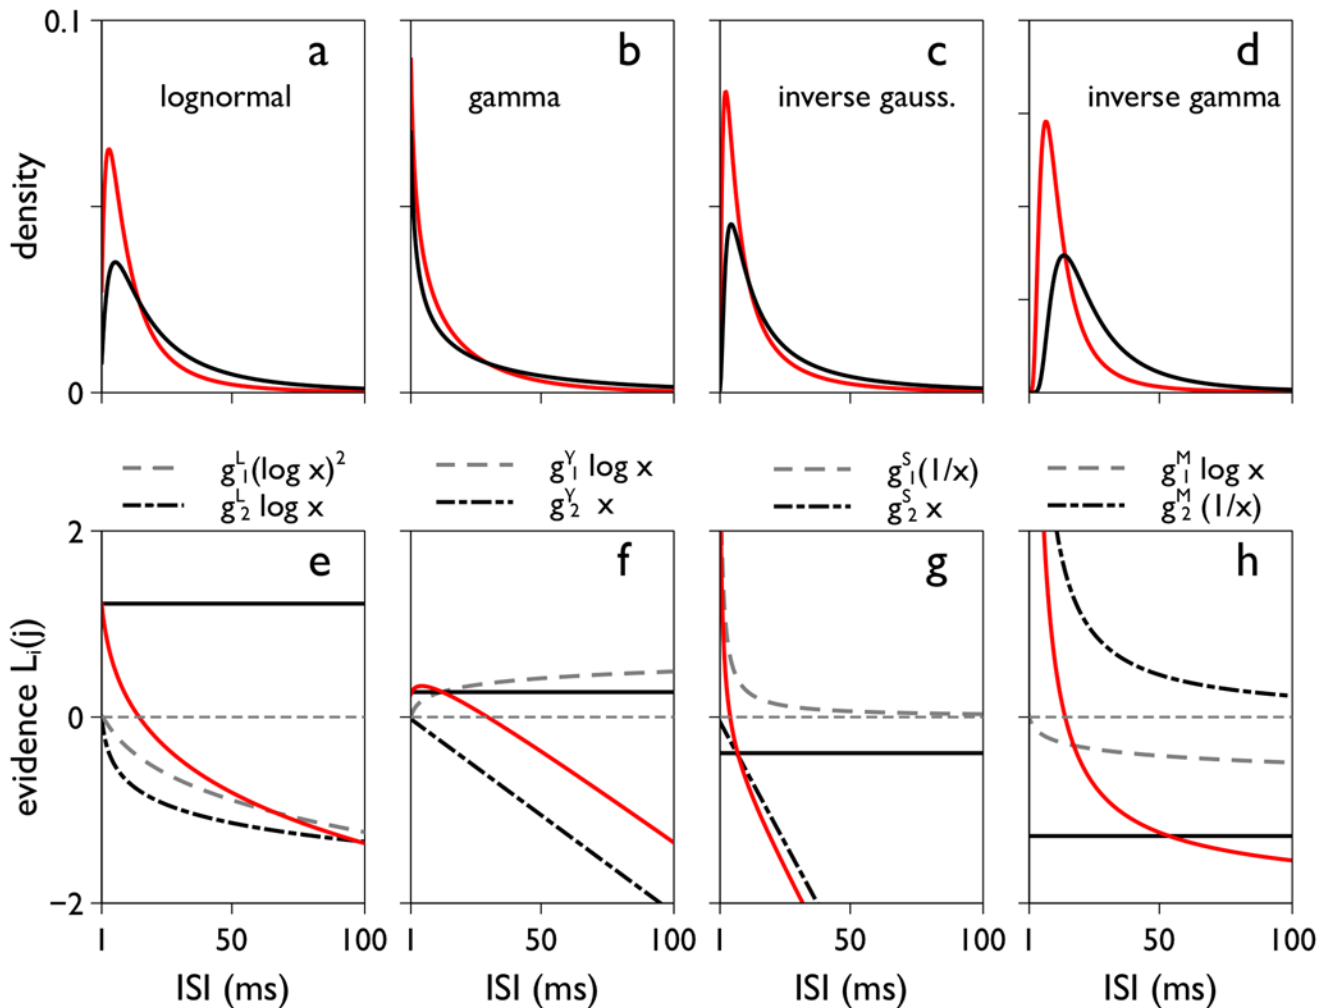
Fig 3. The two-parameter families of pdfs (top) and their 'evidence contributions' $L_i(j)$ (bottom). Panels a-d show the lognormal, gamma, inverse Gaussian, and inverse gamma pdfs respectively, for the independent variance parameter set Ω_{IV} (see Methods). The 'preferred' and 'null' density functions (f_*, f_0) are in red and black respectively. The plots are for ISIs from 1 to 100 ms. For infinitesimal ISIs, the lognormal, inverse Gaussian and inverse gamma tend to zero; for the the gamma the pdf grows up to a bound as the ISI tends to zero. Panels e-f are the corresponding contributions $L_i(j)$ to the accumulated 'evidence' $y_i(T)$ (see Eq 9) and the separate components therein (see Table 1). $L_i(j)$ itself is shown in red, the constant term g_0^D ($D = L, \gamma, S, M$) by the solid black line, and non-constant terms by dashed-grey and broken-black lines. The horizontal dashed grey line indicates 0 on the y-axis.

doi:10.1371/journal.pone.0124787.g003

inverse Gaussian, have an apparently disproportionate effect on $L_i(j)$. However, as noted earlier in connection with the exponential distribution, when summed over spikes, such terms give an expression which is approximately gT (for some gain g). If these terms were identically equal to gT, they may be absorbed into the constant term $B(T)$ in Eq 8, and have a null effect on the posterior. Thus, assuming the terms linear in $x_i(j)$ are a good approximation to gT, they will have a very limited influence on the outcome of a decision.

s-MSPRT has a 'regular' MSPRT counterpart for all ISI distributions

It is instructive to compare the s-MSPRT with a counterpart 'regular' MSPRT with temporally *uniform*-sampling; that is, in contrast to neural spike trains, observations are drawn simultaneously for all channels every time interval δt (MSPRT as in [10, 19, 29]). We denote this

uniform-sampling variant, u-MSPRT. Note this alternative form, like its s-MSPRT counterpart, is not supposed to be based on observations from an underlying continuous process. Rather, u-MSPRT relies on a fundamentally discrete, sequential process with uniform inter-observation time, but is abstract and does not assume any explicit representation of spike arrivals. Its sole purpose is to provide a 'bridge' from the s-MSPRT to the more usual, uniformly sampled decision processes (which may well assume a continuous time foundation). In this scheme, the distribution-independent formalism developed above is preserved almost in its entirety, with $x_k(j)$ interpreted as the j^{th} observation on channel k, drawn from one of the pdfs describing ISIs. Thus, Table 1 applies for all such u-MSPRT, which therefore extends previous results describing the specific form of MSPRT for Gaussian inputs only [10].

In all u-MSPRT variants, the expressions contributing to $y_i(T)$ in Eq 9 now refer to sums up to observation $s(T)$ for any channel (instead of $s_i(T)$) and $T = s(T)\delta t$ ($T_i = T$, for all i). With likelihoods and posteriors thus updated simultaneously for every hypothesis, every δt, u-MSPRT also takes the form in Eq 10 and has the structure in Fig 1.

We now turn to the comparison between s- and u-MSPRT. For both variants, the decision time is governed by the *decision sample*—the number of observations required to reach the threshold. For a given decision, in s-MSPRT the decision sample will depend on the hypothesis which has reached threshold. On average, when $\mu^* < \mu_0$, there will be more observations in a preferred channel than in one of the null channels; this is observed in cortical sites that supply evidence for decision formation, like the middle-temporal visual area (MT) [30] or the primary somatosensory cortex [24]. We report the expected decision sample with respect to an equivalent number of preferred channel observations, $\langle s_s^* \rangle$, irrespective of the hypothesis crossing the threshold. For u-MSPRT, all channels are sampled equally frequently but we will, nevertheless, denote the expected decision sample with respect to the preferred channel, $\langle s_u^* \rangle$, in order to facilitate comparison between u- and s-MSPRT.

For s-MSPRT, the decision time T_D^s is the sum of an integer number of observations, $s_i(T_D^s)$, and a 'residual time' comprising the time from t^0 to first spike arrival at t_i^0. For correct decisions, the expected value of $x_i(0)$ is $0.5\mu^*$. Thus, the expected decision times for each type of MSPRT are given by

$$
\begin{aligned}
\langle T_D^s \rangle &= (\langle s_s^* \rangle + 0.5)\mu^* \\
\langle T_D^u \rangle &= \langle s_u^* \rangle \delta t
\end{aligned}
\tag{15}
$$

Notwithstanding the simple formal relations above, the interpretation of decision making by these algorithms, in terms of an overall *decision time*, is rather subtle. For s-MSPRT, observations are explicitly determined by spike arrivals, and we will ultimately need to know whether we are we dealing with single or multiple spike trains and, if the latter, how these combine to make ISI-pdfs for algorithm input. These questions are taken up again in the Discussion but, in all subsequent results, we consider processing of an individual afferent spike train as the 'unit of decision making'. All the analyses of ISI statistics described above are, therefore, directly applicable. Further, in previous application of uniform-sampled MSPRT to neurobiological decision making, the parameters were chosen to allow *behaviourally* appropriate decision times [10]. Here, however we wish to examine decision making at the level of spike trains and *neural* processing which requires a different approach relating observations in u-MSPRT to ISIs.

Given these issues, we will report results (for both u- *and* s-MSPRT) in terms of the decision sample rather than overall decision times. In particular, we can ask whether the decision samples for each of u- and s-MSPRT are similar. Given the mechanistic difference between s- and u-MSPRT in terms of a homogeneous versus heterogeneous channel sampling rate, it is not clear *a priori* whether this is necessarily the case. However, if they are indeed similar, then

notwithstanding the problems with interpreting decision times, it is instructive to see what choice of sampling period δt would ensure similar decision times. The investigation to answer these questions is largely empirical, but we supply heuristic analysis to give insight into the outcomes.

We ran simulations of s-MSPRT as a function of the number of choices or hypotheses, N, for a range of pdfs, and the two parameter sets Ω_{IV}, Ω_{FV} (see Methods). All the simulation results use trials with an error rate of 5% obtained by iteratively seeking a threshold that satisfied this criterion. Every data point for a particular number of alternatives is the mean over 950 correct, out of 1000 total trials (982 correct for the inverse gamma based s-MSPRT at $N = 2$ with Ω_{FV} in Fig 4 which is at an error rate of 1.8% as it was not possible to reliably achieve a 5% error here; this decision task would appear to be too easy to be compromised to this extent). The large trial numbers ensure that estimation of the error rate during threshold determination is sufficiently accurate. For the simulations in this section, the ISIs are drawn from the distributions, f_* for the preferred channel, and f_0 for the null channels.

The decision samples $\langle s_s^* \rangle$ are estimated in Fig 4 (solid lines). We can consistently interpret these results as decision times because $\mu^* = 16.5$ ms throughout. Then, assuming only the single preferred afferent spike stream, Eq 15 implies that 10 observations translates to 165 ms. The solid symbols in Fig 4 show the comparable decision samples $\langle s_u^* \rangle$ for u-MSPRT. There is clearly a very close correspondence between the two sets of mean decision samples across a range of conditions; that is, $\langle s_u^* \rangle \approx \langle s_s^* \rangle$. There are two possible exceptions to this; the gamma distribution with parameters Ω_{IV} and the exponential with Ω_{FV}, but even here, the correspondence is reasonably good. It would appear therefore that, for a given s-MSPRT, there is a u-MSPRT counterpart with the same mean decision sample and error rate. Further, given the provisos above, Eq 15 implies that decision times in the two cases will be equal (ignoring the residual
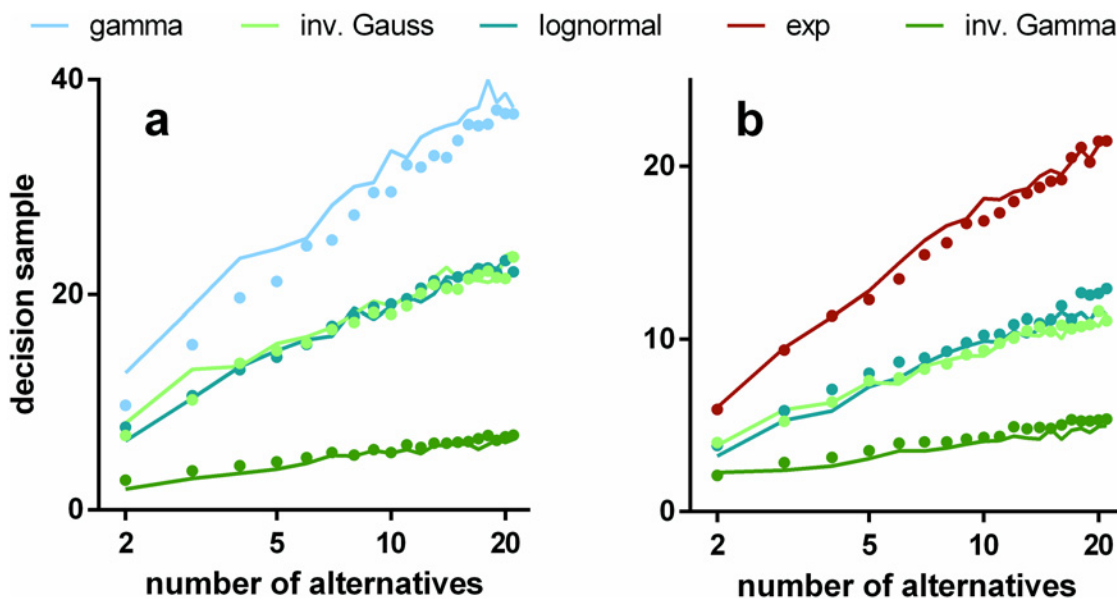

Fig 4. Mean decision samples against number of choices for a range of pdfs, parameter sets, and mechanisms. Each panel shows mean decision sample as a function of the number of choices for a range of pdfs (see legend) and for the two alternative mechanisms: s-MSPRT (solid lines) and u-MSPRT (solid circles). Panels a and b are for the parameter sets Ω_{IV} and Ω_{FV} respectively (see Methods). In the case of Ω_{FV}, the gamma and exponential distributions are identical and so not reported separately. All data points are the mean of 950 correct, out of 1000 total trials (see text for inverse gamma based s-MSPRT). Error bars are omitted for clarity and are small; the standard error of the mean is typically 2% of the mean.

doi:10.1371/journal.pone.0124787.g004

term $0.5\mu^*$) if we assign $\delta t = \mu^*$. In the Methods, we develop a heuristic argument to show why there is a close match between the two methods.

s-MSPRT can be more observation efficient than the usual u-MSPRT

Consider Eq 4. Then, putting $R(T) = \sum_{k=1}^{N} \exp y_k(T)$, Eq 4 becomes

$$P_i(T) = \frac{\exp y_i(T)}{R(T)} \tag{16}$$

where, $R(T)$ is a hypothesis independent normalisation constant. The idea, therefore, is to think of the posterior as a 'scaled' version of $\exp y_i(T)$, although this scaling will change at every spike arrival in any channel. This is illustrated in Fig 2, in which panels b,c therein show $y_i(T)$ and $P_i(T)$ respectively; the notion of scaling is especially clear for the preferred hypothesis. Assuming a small error rate, most decisions will choose this hypothesis, and so the critical feature for the decision time is the trajectory of the integrated evidence therein, $y_*(T)$ (here, and henceforth, a '*' subscript denotes quantities on the preferred hypothesis).

In general, the terms $L_*(j)$ in Eq 9 contributing to $y_*(T)$ have similar form whether they are notionally obtained from spike arrivals or uniform sampling. Therefore, at the level of observations, typical trajectories for $P_*(T)$ will be similar in u- and s-MSPRT, up to an overall scaling by $R(T)$. However, even with identically shaped trajectories, the decision sample will depend on $R(T)$ and the threshold in each case. We therefore proceed to examine these quantities.

In evaluating $R(T)$, we first we note that each contributory term $L_i(j)$ to the evidence in the null hypothesis (in either MSPRT) is likely to be negative, since we expect $f_*(x_i(j))/f_0(x_i(j)) < 1$. This, in turn tends to make $y_i(T) < 0$; as an example of this, see Fig 2b. Secondly, suppose we have taken the same number of observations in both u- and s-MSPRT on the preferred channel. Then, since $\mu^* < \mu_0$, there will be fewer observations in a null channel for s-MSPRT than for u-MSPRT because, for the former, they are sampled μ_0/μ^* times more slowly than those on the preferred channel, whereas all channels are sampled at the same rate for u-MSPRT. Incidentally, note that given the equality of decision samples (with respect to the preferred channel) this means that the total number of scalar observations, summed across all channels, to reach decision is less for s-MSPRT than it is for u-MSPRT. Thence, s-MSPRT is more *observation efficient* than its non-spiking counterpart.

There is no single optimal ISI distribution for decision making

There is a clear distinction in Fig 4 between the performance of algorithms assuming the different distributions. Is the rank ordering of performance maintained as we vary the distribution statistics? To explore this we repeated the experiments with s-MSPRT corresponding to $N = 10$ in Fig 4, but with other parameter sets $\hat{\Omega}_{IV}(\mu_0), \hat{\Omega}_{FV}(\mu_0), \mu_0 = 49.5, 66, 82.5$ derived from the original sets Ω_{IV}, Ω_{FV} (see Methods). The results are shown in Fig 5, which also show those for Ω_{IV}, Ω_{FV} for comparison ($\mu_0 = 33$ ms).

Fig 5a shows a clear difference in relative performance over the parameter sets $\Omega_{IV}, \hat{\Omega}_{IV}(\mu_0)$ defined by their means. Further, there is some patterning to the variation in which the peak decision sample is a unimodal function of mean ISI. The inverse gamma shows best performance for small means, but is superseded by the other distributions at higher means, and in particular, by the gamma distribution. However, the latter is not a good fit to the data sets we examined (see Fig 10) and so a realistic mechanism might be more likely to use the lognormal or inverse Gaussian, both of which have similar decision samples.
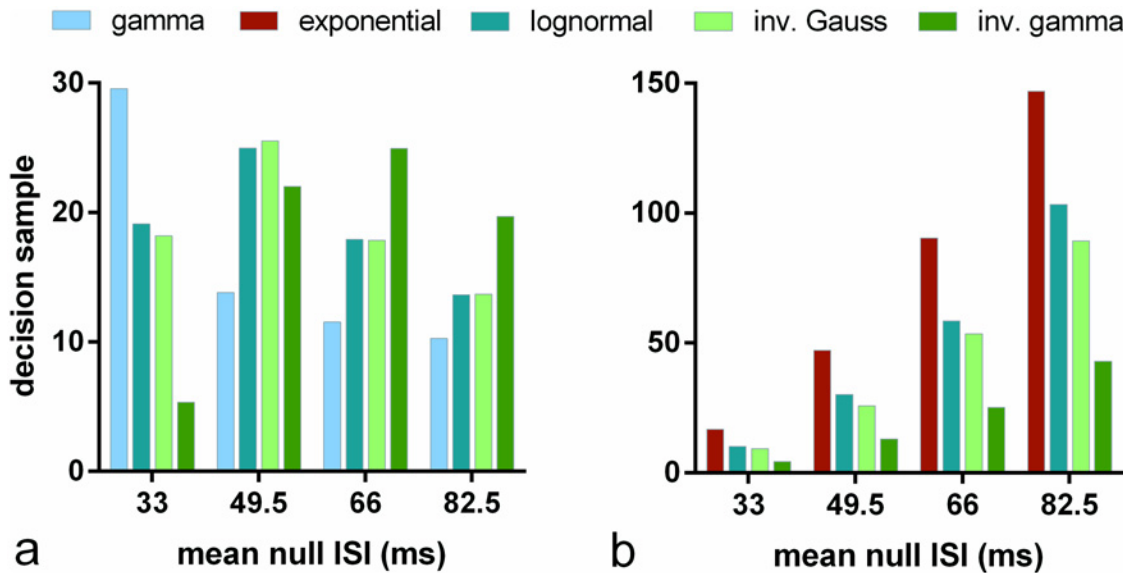
Fig 5. Decision sample in s-MSPRT against mean ISI for range of pdfs and parameter sets. Each bar shows, for the pdf indicated in the legend, the mean decision sample for N = 10 alternatives, averaged over 950 correct out of 1000 total trials. Panel a used parameter sets $\Omega_{IV}, \hat{\Omega}_{IV}(49.5), \hat{\Omega}_{IV}(66)$, $\hat{\Omega}_{IV}(82.5)$, panel b used $\Omega_{FV}, \hat{\Omega}_{FV}(49.5), \hat{\Omega}_{FV}(66), \hat{\Omega}_{FV}(82.5)$ (see Methods). Each group of bars relates to one parameter set with its μ_0 indicated on the x − axis (Ω_{IV}, Ω_{FV} have $\mu_0 = 33$). For the case of Ω_{FV} and any $\hat{\Omega}_{FV}(\mu_0)$, the gamma and exponential distributions are identical and so not reported separately.

doi:10.1371/journal.pone.0124787.g005

In contrast, Fig 5b shows no such variation of performance for the sets $\Omega_{FV}, \hat{\Omega}_{FV}(\mu_0)$, defined by their means and variances; the rank order is preserved. The inverse gamma also shows best performance with this parameterisation. In sum, there would appear to be no consistently 'best' distribution for decision making and that distribution contingent performance varies with the statistics of the data distributions.

Note also that the decision samples under the parameter sets $\hat{\Omega}_{FV}(\mu_0)$ in Fig 5b are much larger in general than they are for those derived from Ω_{IV} in Fig 5a. This is because, for $\hat{\Omega}_{IV}(\mu_0)$, the standard deviation for f_* and f_0 does not grow with μ_0, whereas it does for $\hat{\Omega}_{FV}(\mu_0)$, since it is always equal to μ^*, μ_0.

MSPRT is robust under variation of hypothetical distribution

Thus far we have assumed that the underlying statistics of the spike trains are the same as those of the hypothetical distributions (those in the likelihood). In particular, we have assumed that the functional form of the spike ISI distributions f_*^s, f_0^s were identical to their hypothetical counterparts in the decision mechanism, f_*^h, f_0^h, (so that no distinguishing superscript h, s was required). However, in general, the decision mechanism may not 'know' *a priori* what form f_*^s, f_0^s take. We now ask the question: what effect would an incorrect choice of the pdf form f_*^h, f_0^h take on the decision time? To investigate this we fixed the distributions of ISIs to be inverse gamma, and supplied this data to decision mechanisms (u-MSPRT) based on a variety of pdfs. We used the Ω_{IV} parameter set for both distribution sets throughout. The results are shown in Fig 6 by the pattern filled bars. There is always a change in performance when an 'incorrect' hypothetical distribution form is used. However, for all the 'incorrectly' established mechanisms, the performance is better with the inverse-gamma-distributed data than that when each mechanism
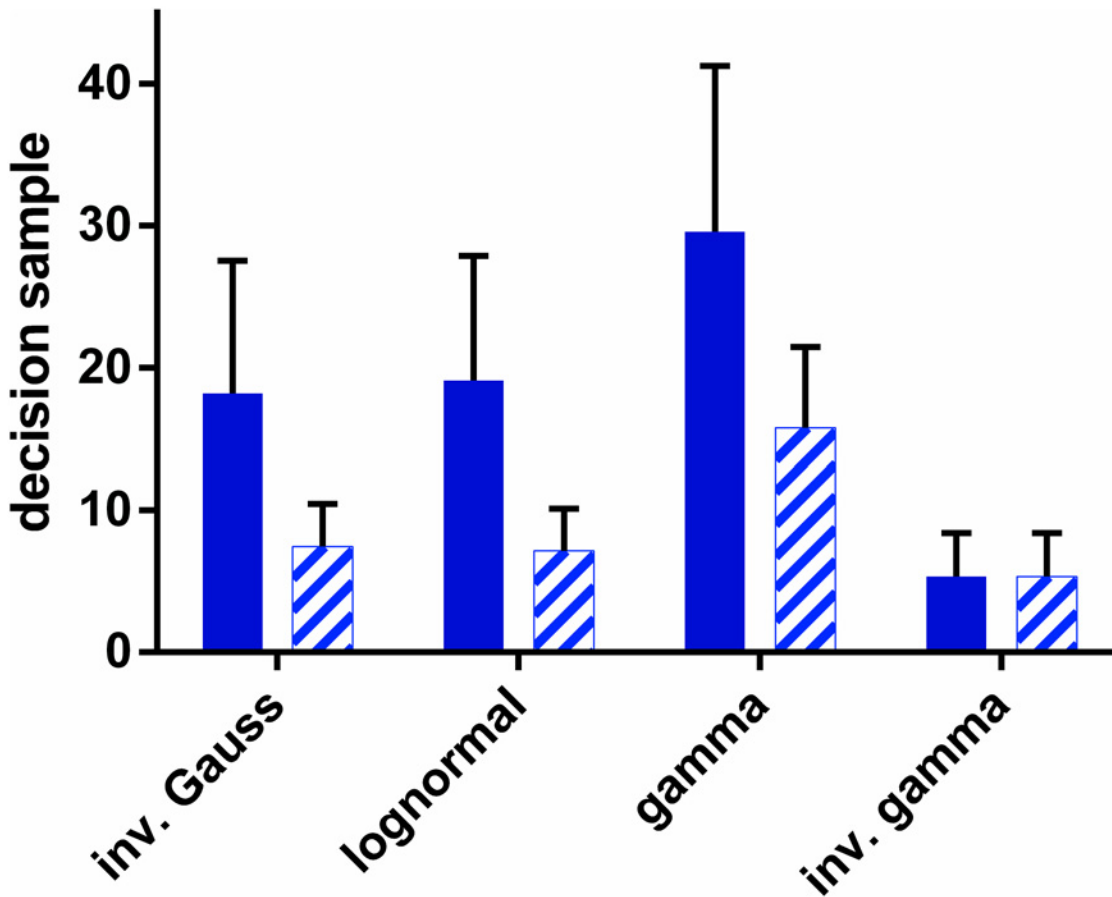
Fig 6. Decision sample for u-MSPRT when the distributions inthe data were not matched to those tested for. Each bar shows the mean decision sample, for $N = 10$ alternatives, averaged over 950 correct out of 1000 total trials. The parameter set was Ω_{IV}. The pale, patterned bars are for the case when the data is always sampled from an inverse gamma distribution, but inserted into mechanisms which test using the distribution indicated on the x – axis (by definition, the bars have equal height for the inverse gamma). The solid bars are for the case when the tested-for distribution matches the true distribution of ISIs, as indicated on the x – axis. Error bars are at one standard deviation.

doi:10.1371/journal.pone.0124787.g006

uses observations drawn from pdfs that are the same in the data and the hypothesis test (solid bars in Fig 6). Thus, there is no catastrophic decline in performance when using non-matching hypothetical distributions, and performance variation appears more intimately linked to the characteristics of the observations themselves.

Expected total information gain for a decision is constant

It is clear from several of our results that the choice of pdf can make a substantial difference to the decision making performance. However, it is not clear what characteristics of the choice of pdf cause these differences. We might suppose that performance will be a function of how 'far apart' are the distributions f_*, f_0 and one popular measure of this distance is the Kullback-Leibler divergence (KLD) between two pdfs $p(x)$, $q(x)$ [33], defined by

$$D(p||q) = \int_{-\infty}^{\infty} p(x) \log_2 \frac{p(x)}{q(x)} dx \qquad (17)$$

Here we will use base-2 logarithms, so that results can be reported in bits. Note the KLD is, in general, asymmetric with $D(p\|q) \neq D(q\|p)$ (although symmetry may occur under some circumstances—*e.g.* Gaussians with different means and the same variance).

Now let $p \equiv f_*$, $q \equiv f_0$, then, taking expected values of the per-observation ratios in Eq 9 with respect to f_*

$$D(f_*\|f_0) = \langle L_i(j)\rangle_{f_*} \qquad \forall i,j \tag{18}$$

The KLD is therefore the mean increase in log-likelihood, $y_i(T)$, per observation [41]. Denoting quantities on the preferred channel by *, and using Eq 9 together with Wald's identity [42, 43], we can find the corresponding expectation of the accumulated evidence for the preferred hypothesis at decision time T_D^s,

$$\langle y_*(T_D^s)\rangle_{f_*} = \langle s_s^*\rangle \langle L_*(j)\rangle_{f_*} \tag{19}$$

Thus, using Eq 18

$$\langle s_s^*\rangle = \frac{\langle y_*(T_D^s)\rangle_{f_*}}{D(f_*\|f_0)} \tag{20}$$

We now turn to an empirical investigation of Eq 20. If the numerator is constant for variation in parameters or pdf, we expect a simple inverse relation between $\langle s_s^*\rangle$ and $D(f_*\|f_0)$. Fig 7a and 7b show that this is indeed true for each of the pdf classes Ω_{IV}, Ω_{FV} ($\mu_0 = 33$ ms), used in Fig 5a and 5b respectively. The dashed lines are power law fits and are remarkably good: $R^2 > 0.998$, with exponents -0.866, -0.844, for Ω_{IV}, Ω_{FV} respectively. These exponents are almost -1, as predicted by Eq 20 and, it would therefore appear that KLD is a good predictor of 'local' variations in performance under a given parameter set. However, plotting $\langle s_s^*\rangle$ against $D(f_*\|f_0)$ for all 32 tests in Fig 5 indicates a more general result—see Fig 7c, where the log axes emphasise the power law fit over the wide range of the variables. The fitted function is

$$\langle s_s^*\rangle = 4.594 D(f_*\|f_0)^{-1.034} \tag{21}$$

and again, the fit is very good ($R^2 = 0.997$). Just as significantly, the exponent is very close to -1, in which case the product of KLD and the decision sample is a constant A

$$\langle s_s^*\rangle D(f_*\|f_0) = A \tag{22}$$

Some indication of why this might be true is supplied by Veeravalli and Baum [44] (see Methods). This indicates that the 'constant' A depends only on the error rate ϵ and number of hypotheses N according to

$$A(\epsilon, N) = \left(1 - \epsilon - \frac{\epsilon}{N-1}\right) \log \frac{(1-\epsilon)(N-1)}{\epsilon} \tag{23}$$

(natural logarithm, note that in Fig 5, ϵ and N are fixed).

Further, Eq 22 also has theoretical plausibility grounded in the notion that KLD is a foundation of information theory. In fact, in its original formulation, it was noted that the KLD gave the mean *information per observation* for discrimination between two hypotheses about the distribution leading to the sample [33]. This kind of interpretation leads to the KLD also being referred to as an *information gain*. To emphasise this view we write $D(f_*\|f_0) \equiv I(f_*, f_0)$. Then, using this, and Eq 23, the empirically supported result in Eq 22 becomes

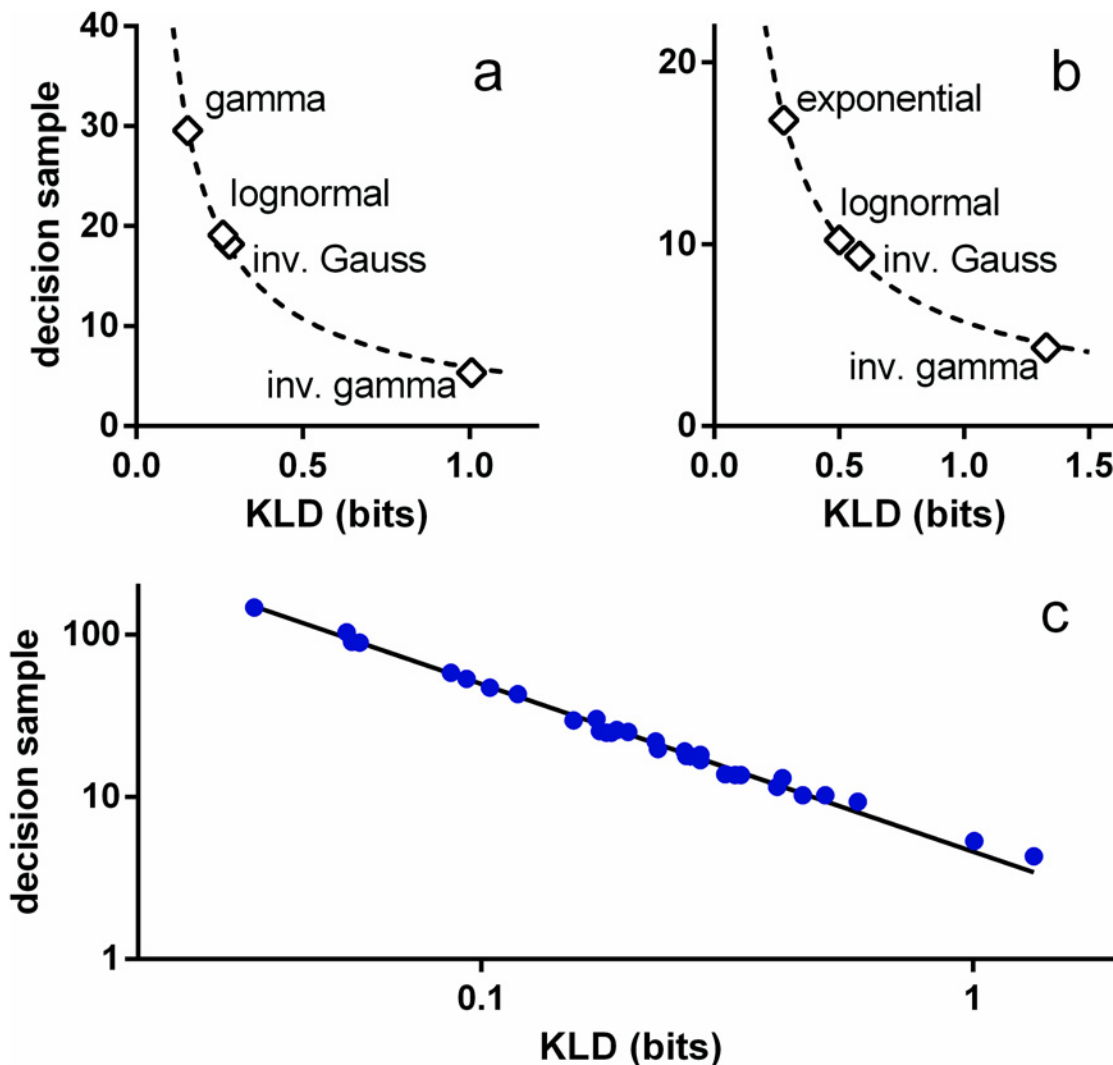$$\langle s_s^*\rangle I(f_*, f_0) = A(\epsilon, N) \tag{24}$$

Fig 7. Mean decision sample for s-MSPRT against KLD. The KLD in all cases is $D(f_*\|f_0)$ (see text). Panels a, b are for decision with the parameter sets Ω_{IV}, Ω_{FV}, respectively, and both use $N = 10$. The data points are the open symbols and the dashed lines, best fit power laws (nonlinear least squares). In panel c, the data points shown in blue symbols correspond to all the decision samples in Fig 5. The solid line is the best fit power law.

The left hand side of Eq 24 is the product of the mean information gain per observation and the expected number of observations to decision; that is the mean total information required for the decision. Thus, Eq 24 states that, for a given N and ϵ, the mean total information for a decision is constant. In this view, the empirically determined constant $A = 4.594$ bits (Eq 21) is the expected amount of information required to perform a decision among 10 alternatives, with a 5% error rate, given uncertainty in the signals supplied. We now go on to use the result on information conservation to show how the neural refractory period can facilitate decision making and to explain two well known psychophysical phenomena.

The spike refractory period is a benefit, not an impediment, for decision making. In real biological neurons, spikes occupy a finite width and there is a refractory period between spikes which forces a lower limit on the ISI of not less than 1 ms. A popular choice for spike

interval generation is the Poisson process for which the ISIs are distributed according to an exponential form $f_e(x)$. However, $f_e(x)$ has its mode (maximum) at $x = x_{mode} = 0$ thereby allowing the unrealistic occurrence of arbitrarily small ISIs. Nevertheless, if the assumption of $f_e(x)$ demonstrated a decision making procedure (*e.g.* s-MSPRT) with a performance advantage compared to other distributions, it could be that real neurons work to ensure their ISI distributions are as close to $f_e(x)$ as possible. On the other hand, if assuming $f_e(x)$ gives inferior performance compared to distributions for which $x_{mode} > 0$, the ostensible limitation in neural processing that is the refractory period may be thought of, instead, as a enhancing feature, since it is this mechanism that has forced a non-zero location of the mode.

To cast light on this issue, consider the results in Fig 4b. This shows a comparison of the performance of s- and u-MSPRT assuming f_e against a range of other distributions with $x_{mode} > 0$ using parameter set Ω_{FV}. The procedures assuming the exponential distribution clearly give longer decision times compared to the others. This could be a peculiarity of the choice of the means, but examination of Fig 5b shows otherwise; in all cases the rank ordering of performance across distributions is maintained and the exponential performs worst. Similarly, the gamma assuming (u-)s-MSPRT performs worst among its cohort under Ω_{IV} (Fig 5b) while assuming distributions with maximum density at infinitesimal values of x (Fig 3b). By Eq 17, the KLD is determined by the ratio of the entirety of a pair of densities. When the modes of both are greater than 0 they also tend to differ from one another, as it happens in general in neural ISIs under preferred and null conditions, as well as in the rest of our densities under biologically realistic parameters (Fig 3a, 3c and 3d). Then, the area under the log-ratio of such densities and thus the KLD between them tend to increase which in turn, by Eq 24, improves the expected performance of a decision procedure. All this therefore supports the hypothesis that the existence of a refractory period helps neural decision making by facilitating the contrast of the ISI distributions of simultaneous spike trains.

Information conservation explains Hick's Law. Hick's law (sometimes known as the Hick-Hyman law) describes the relation often observed between mean reaction time T_R and the number of alternatives N in choice reaction time experiments with low error rates [35, 45, 46]. There are two slightly different variants

$$
\begin{aligned}
T_R &= a \log (N + 1) \\
T_R &= T_0 + a \log (N)
\end{aligned}
$$

(25)

where the first form was that originally proposed by Hick and the second used in more recent applications [47]. The second form is also more plausible as the term T_0 can absorb irreducible components of reaction time which originate in some minimum decision time, sensory processing delay and motor execution. In fact, in tasks where the stimulus-response mapping is too simple or over-learned (*e.g.* [48, 49]), the contribution of the logarithmic, N dependent factor appears to become undetectable [8].

The use of a log axis for N in Fig 4 emphasises the fit of the simulation data to laws of this form (all linear regressions of both kinds in Eq 25 have residuals $R^2 > 0.95$). Our results are therefore consistent with the psychophysics of choice experiments as expressed in Hick's law.

Moreover, the information gain perspective is able to explain the general form of Hicks law. Starting from Eq 23, and assuming $\epsilon/(N-1)$ to be negligible (due to small ϵ)

$$
\begin{aligned}
\langle s_s^* \rangle I(f_*, f_0) &\approx (1 - \epsilon)\left(\log \frac{(1-\epsilon)}{\epsilon} + \log(N-1) \right) \\
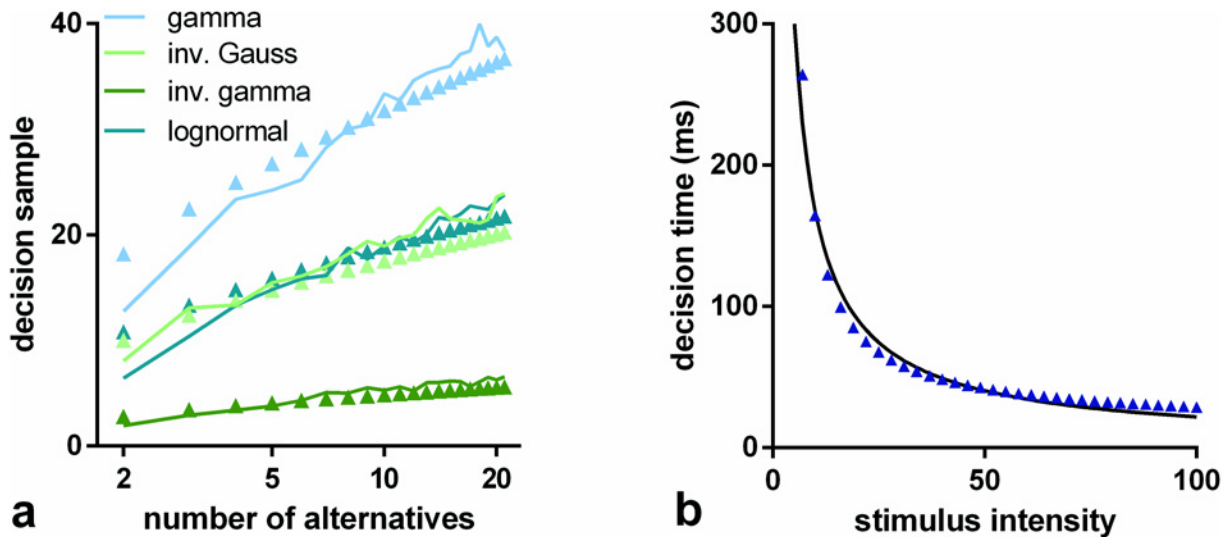&= h(\epsilon) + (1 - \epsilon) \log(N - 1)
\end{aligned}
$$

(26)

Fig 8. Hick's and Piéron's laws from conservation of information. Panel a is a direct counterpart of Fig 4a. The decision samples for s-MSPRT are shown as solid lines and the predictions from Eq 27 shown as solid symbols. Panel b shows results of the virtual experiment derived from Eq 30 (blue symbols) and a best fit power law (solid black line).

with $h(\epsilon) = (1 - \epsilon)\log[(1 - \epsilon)/\epsilon]$. For given neural distributions, $I(f_*, f_0)$ is constant and so

$$\langle s_s^* \rangle = a + b \log(N - 1) \tag{27}$$

which is close to the second form of Hicks law (in terms of a decision sample) given in Eq 25. Eq 27 is similar to the empirical expression for mean decision time of u-MSPRT obtained by McMillen & Holmes [8]. Through a different method, our expression completes their intuition, confirms the $\log(N - 1)$ dependence on N and demonstrates this form to generalize to s-MSPRT; both expressions constitute concrete, experimentally testable predictions.

Further, we tested the ability of this theoretical result to quantitatively explain our data, by calculating $I(f_*, f_0)$ for the distribution pairs used in Fig 4a, and substituted these in Eq 26; the results, together with the original simulation data are shown in Fig 8a There is good agreement with the empirical results.

Information conservation explains Piéron's Law. There is another widely observed, lawful relationship in psychophysical experiments which has a long history. Piéron's law [36, 50] is a psychophysical regularity where the simple reaction time to the detection of stimuli across a range of sensory modalities has been found to depend through a power law on the intensity of the stimulus, u. It is also a characteristic of low ($\lesssim$ 5%) error rate circumstances [51, 52]. Piéron's law has also been extended to the effect of stimulus intensity on choice reaction times (CRTs) [51, 53, 54] and to the effect of stimulus separability on CRTs [55]. In terms of the influence of stimulus intensity u on reaction time T_R, Piéron's law may be written as

$$T_R = T_0 + au^{-b}, \qquad b > 0 \tag{28}$$

Just as in Eq 25, the term T_0 may include components of reaction time which originate in minimum decision time, peripheral sensory delays and motor execution; we are interested only in the main decision making component au^{-b}.

To make the link between this component of Piéron's law and Eq 24 we need a *linking hypothesis* which describes how neural responses, as firing rates, depend on u; that is we require firing rate r as a function of u. We base this on the study by Muniak et al on firing rates and tactile sensation [56]. Here, the task is one of stimulus detection and, while we are interested primarily in choices tasks, we take the kind of linking hypothesis developed in that study as indicative of a more general one. Further, we modified the relation $r(u)$ used by Muniak et al [56] to include an explicit baseline firing rate $r(u_\theta)$, where u_θ is the sensory threshold.

$$r(u) = \alpha \log \frac{u}{u_\theta} + r(u_\theta) \tag{29}$$

The use of a baseline ensures that we are comparing the firing in the preferred channel with a non-zero null rate; that is we assume $\mu_0 = 1/r(u_\theta)$ for a fixed f_0. Similarly, we assume the pdf of ISIs in the preferred channel, $f_*(u)$, is parameterised by its mean ISI $\mu^*(u) = 1/r(u)$.

Then, starting from the decision time relation, $T_D(u) = \langle s_s^* \rangle \mu^*(u)$, and using the expression for $\langle s_s^* \rangle$ from Eq 24, we have

$$T_D(u) = \frac{A(\epsilon, N)}{\left[\alpha \log \dfrac{u}{u_\theta} + r(u_\theta) \right] D(f_*(u)\|f_0)} \tag{30}$$

To test the validity of Eq 30 we ran a 'virtual experiment', obtaining simulated data points from Eq 30 with the following parameterisation, chosen to give firing rates comparable to those in the rest of this article: $\alpha = 10$, $r(u_\theta) = 10$, $u_\theta = 1$, $7 \leq u \leq 100$; f_*, f_0 were lognormal with, $\sigma_0 = 200$ ms, $\sigma^* = 65$ ms; the error rate $\epsilon = 0.05$, and $N = 2$. Fig 8b shows the resulting theoretical datapoints (dark blue symbols) and a best fit line of the form $T_D(u) = au^{-b}$. The fit is good, and so it is possible for virtual experiments based on Eq 30—derived in turn from Eq 24 —to be described by a Piéron-like law. We infer that Eq 24 with a linking relation like Eq 29 could account for Piéron's law *in vivo*.

Discussion

We have shown how arbitrary renewal spike trains may be subject to general Bayes-based sequential analysis and, when this treatment specialises to the use of log-probability ratios, we obtain an instance of the MSPRT. In this spiking MSPRT (s-MSPRT), the data are ISIs, and the sampling times are given by the asynchronous spike arrival times on several parallel channels. The corollary of this is that the posterior, for any hypothesis, is updated at spike arrivals on any channel.

Our analysis of ISI data for neurons involved in a decision task showed that this is not well approximated by a Gaussian. Indeed there is a conceptual problem that the Gaussian admits negative ISIs which is infeasible. On application of s-MSPRT to the most oft-used, realistic ISI distribution –the exponential yielding 'Poisson spikes'– the accumulated evidence, $y_i(T)$, up to time T, is shown to be approximately equal to the scaled count of spikes to T (up to a scaling factor). This grounds previous work by Zhang and Bogacz (2010) [17] whose starting point was the *assumption* that the evidence in this case was given by a (scaled) spike-count. Our result also modifies the result of [17] showing that, while providing the main source of evidence, for an exponential-based s-MSPRT the spike count needs to be augmented by an additional 'correction' ($\delta_i(T)$, in Eq 12).

Our data analysis also highlighted a series of candidate distributions which fitted the data better than the exponential (gamma, inverse gamma, lognormal, inverse Gaussian). The

accumulated evidence for these was shown to be a spike count with additional terms which, in general, cannot be neglected.

We showed that s-MSPRT is comparable with a uniform sampling counterpart—u-MSPRT—if the sampling rate of the latter is set to the mean ISI on the preferred channel, μ^*. However, in making this comparison, we argued that performance is best measured in most cases using the decision sample rather than decision time *per se* as the latter requires careful interpretation (see below). The decision samples for s- and u-MSPRT are very close and there is a heuristic argument to be made why this might be so. Nevertheless, s-MSPRT can achieve the same performance and error rate with fewer observations than u-MSPRT because it doesn't have to accommodate as many observations from the non-preferred ('null') channels; hence, s-MSPRT is more observation efficient than u-MSPRT. As with many decision mechanisms (*e.g* [8, 10, 57, 58]), both s-MSPRT and u-MSPRT showed a Hick's law relation between decision sample and number of alternatives.

A key result was that there is no universally optimal distribution for making a decision. The best performance varies with the underlying statistics of the data. However, a decision mechanism assuming the exponential distribution shows, in general, an inferior performance compared to when assuming other distributions with mode at non-zero ISI. Thus, the spiking refractory period can be viewed as positive feature which aids decision making mechanisms, rather than a drawback which hinders attainment of optimal performance therein.

In most of our simulations, the s-MPSRT was configured to give the best performance because the hypothetical distributions (*i.e.* those tested for) were matched to those of the data. That is, the decision mechanism was given 'privileged' information about the data. However, this will not normally be the case and we showed a case where this assumption is violated and there is no catastrophic decline in performance.

We showed empirically in our model that the mean total information gain required to make a decision was constant, over a wide range of distributions and parameters. This was supported by some analytic results and so we conjecture that it may be a quite general result. While this result may not sound surprising in isolation, it is remarkable that it emerges empirically in such a precise way from the examination of a variety of conditions with different distributions and parameterisations. This result allowed us to provide a theoretical rationale for Hick's law and Piéron's law and a concrete, experimentally testable expression for mean decision time as a function of sensory information content, error rate and number of choices (Eq 24).

Interpretation of s-MSPRT and perceptual decision times

The grounding of a probabilistic decision mechanism in neural firing statistics is clearly a conceptual advantage, since neural decision mechanisms must ultimately make use of spikes in some way. However, we are confronted with interpretational issues which may not arise if we are working at a more abstract level of description, in which 'evidence' and 'sampling times' may be chosen to fit behavioural reaction times. First, assuming the brain performs decisions using a mechanism akin to s-MSPRT, we suppose it operates in neural systems which do not include sensory pre-processing (through early visual cortex, for example) or delivery of motor output (through motor cortex to spinal cord or brainstem). In reality, most probably these three stages are not sequential but overlap in time [59]; indeed, decision making may carry on during motor execution resulting in observable 'changes of mind' [60]. Nevertheless it is useful to conceive of these three contributions as separate processes as we are interested in the central, decision making stage, which will be observable electrophysiologically in studies like [30, 61–67].

Given this demarcation of processing, there is a wide range of possible neural decision times discernible in perceptual tasks. For example, using high contrast, easily distinguished stimuli in a paradigm designed specifically to tease out the perceptual decision making process, Stanford et al [68] showed that this could occur in as little as 30 ms for a 2-alternative task. However, for the relatively hard to distinguish stimuli in the random dot motion task (RDMT) with 2–4 alternatives, neural integration times can be several hundred ms [69, 70].

The statistics of the neural data we used here were obtained in a RDMT with stimulus coherence of 12.8% [30] (see Methods). In the RDMT-based study of Churchland et al (2008) [69], with stimulus coherence of 12.8%, the mean reaction times for $N = 2$ and $N = 4$ were 535 ms and 618 ms respectively. Using estimates of sensory and motor delays of 200–300 ms [60, 71] we would expect corresponding decision times of 285 ± 50ms, and 368 ± 50ms; that is of the order of 10^2 ms. A full behavioural comparison with studies like this is out of the scope of the present article, however, if we convert the sampling times for small N in Fig 4a to decision times using Eq 15, we obtain decision times of the same order as this for several of the distributions, including the lognormal which gave the best fit to the data in [30, 31].

Thus far, we have predominantly used a simple version of s-MSPRT based on a single neural input stream on each channel (Fig 1b). In general, however, the theory (leading to Eq 10) is neutral regarding the origin of the spikes in each channel—they may derive from a single or multiple neural sources (Fig 1a). We now consider two possibilities for combining the information across them. In the first, we assume there is some mechanism of *spike superposition* to pool spikes across neurons into a unitary stream before being subject to s-MSPRT. In this case, for large numbers (~ 100) of contributing streams with arbitrary distributions, the resulting superposition of spike trains has a distribution extremely close to the exponential (albeit with some subtleties of its power spectrum) [72]. If μ^*, μ_0 are the mean ISIs within each input stream for preferred and null channels respectively, then the *effective* means of the ISIs in the superposed streams in each channel for M inputs per channel are μ^*/M, μ_0/M. However, assuming a good approximation to the exponential distribution, the gain, $g_0^E = \log \lambda_*/\lambda_0$, remains unchanged, and so the accumulated evidence $y_i(T)$ is still a spike count which simply scales with M. This will render shorter decision times and lower error rates, as more information is available.

The second approach works by directly pooling observations over M inputs. It is obtained by directly extending the formalism leading to s-MSPRT to include more inputs than hypotheses ($C > N$). Thus, let $x_k^r(j)$ be the ISI from the r^{th} spike train of channel set k at time $t_{k,r}^j$, and let $T_k^r = \max_j\{t_{k,r}^j : t_{k,r}^j \leq T\}$, then it is straightforward to show that the analogue of Eq 9 is

$$y_i(T) = \sum_{r=1}^{M}\sum_{j=1}^{T_i^r} L_i^r(j), \qquad L_i^r(j) = \log \frac{f_*(x_i^r(j))}{f_0(x_i^r(j))} + B(T) \qquad (31)$$

Mechanistically, since the order of summation can be reversed, the operations of integrating likelihood ratio terms over time within a stream, and summing across contributory streams, can be conducted in either order. Since the individual ISI statistics are preserved in each stream, the individual $L_i^r(j)$ are distributed in the same way as their single stream ($M = 1$) counterparts. Using a slight abuse of the original terminology we refer to this method as *evidence oversampling*, as it is equivalent to a version of the original s-MSPRT which samples M-times faster than the rate given by the nominal ISI sampling times. We would therefore expect its decision times to scale, on average, as $1/M$.

For large M, both spike superposition and evidence oversampling may comprise *dense coding* implementations (that is, actively engage a large fraction of the relevant neural population). *Sparse coding* ones could be devised by adding: (a) input spike trains that are weakly- or non-informative for the decision and (b) corresponding distributions in the likelihood functions which make weak contributions to them. All approaches in this study assume that spikes, and functions of spike statistics like $L_i^r(j)$, are independent over different inputs. This is almost certainly not the case. Studying the net effect of correlations on decision making remains an interesting problem since positive correlations among spike trains [73, 74] may be detrimental for discriminability, but negative serial ISI correlations may reduce the variability of the signal [75] and help on information transfer [76]. However, our simple independent-input cases, provide a starting point for future models of pooling evidence in these ways. Whatever the details of any ensuing models, the range of single neuron ISI statistics which can yield physiologically plausible decision times will be extended considerably by allowing combination or pooling over many inputs ($M > 1$).

Optimality or sufficiency?

It is often assumed that the brain implements nearly-optimal decision making mechanisms. For $N = 2$, u-MSPRT collapses [29] to Wald's sequential probability ratio test [77]. u-MSPRT is thence optimal in the sense that it minimizes the mean sample size (and decision time) to make decisions at any given error rate; for $N > 2$, it is asymptotically optimal as it minimizes it for vanishingly small error rates [10, 29].

Beck et al (2012) [32] have recently shown that decision mechanisms will, in general, be implemented in a fundamentally suboptimal way. The main reason they give for this is that the brain does not usually have access to the statistical structure of the task. This was precisely the issue addressed in the Results dealing with robustness. There, the true statistics of the stimuli were supplied by an inverse gamma distribution, but we used incorrect (non-inverse-gamma) hypothetical pdfs for the computation of probability ratios. This mismatch of hypothetical and real, underlying data distributions is inevitable; the real data will never comply exactly with an analytically tractable pdf. Such a mismatch may also occur when the response distribution of sensory neurons changes by learning or adaptation (as in [78–81]) and stops approximating 'previously trained' hypothetical distributions (although this seems not to happen in MT over RDMT training [82]).

The results in our experiment were that low decision samples were maintained and were, in all cases, lower than those of other mechanisms assuming non-inverse-gamma pdfs used in conjunction with their own, correctly matched data. In this case, therefore a deterministically sub-optimal decision mechanism (in the sense of Beck et al) *suffices*, because its potentially optimised variant (here with inverse gamma data matching the test-pdf) is better than its counterparts.

There is a related argument here which starts with the data. The dataset we used to establish the parameters Ω_{IV} was best fit by the lognormal, but MSPRT assuming this distribution does not have the best performance; this place is occupied by that assuming the inverse gamma. Thus, if the best performance (given the set of pdfs we used) requires data with an inverse gamma distribution (and ideally, inverse gamma hypothetical pdfs) why is the data not distributed in this way? We argue that, in general, neural codes will putatively result from a compromise in achieving several goals, including maximising discriminability of firing patterns, as discussed here, and expressing an efficient transformation of their inputs based on perceptual information [83–85].

Neural substrates for s-MSPRT

While the s-MSPRT is grounded in the neural *communication* medium of action potentials (spikes) we have yet to address the way in which specific *computations* involved in mechanisms like s-MSPRT may be performed. Consider first the complete Bayes-based expression in Eq 4. This apparently complex form is subject to a mapping onto the basal ganglia—a subcortical group of interconnected nuclei involved in mediating action selection. Thus, Bogacz and Gurney (2007) [10] have shown that (at least simplified forms of) the basal ganglia architecture and cortex could implement Eq 4 when applied to a u-MSPRT, using inputs of Gaussian form. While the original MSPRT was not confined to Gaussian inputs [29], its expression in the form given in Eq 4 is essential for the mapping to basal ganglia-based architectures. Lepora and Gurney [19] have since extended this mapping to much wider range of decision processes (with arbitrary type, and numbers of, pdfs). Indeed there are several possible mappings from this form to basal ganglia and associated circuitry [10, 19, 86]. One possible substrate for a process like the s-MSPRT as a whole therefore, is the basal ganglia and its afferent and target structures.

In all of the anatomical mappings of MSPRT to basal ganglia, the cortex is the locus of the neural representation of integrated evidence $y_i(T)$. Moreover, it has often been assumed that the pdfs delivering input to these integrators are Gaussian (although this is not necessary [17, 19, 29, 37–39, 87]). The integration of Gaussian signals is straightforward if all Gaussian inputs have the same variance; it is simply the product of a gain and a term linear in the (abstract) signal input $L_i(j) = x_i(j)$ [7, 10, 16]. However, the forms for $y_i(T)$ resulting from the distributions we consider here can be complex (see Table 1). We now offer some heuristic arguments to suggest how the accumulated evidence $y_i(T)$ might be computed in cortex. In order to ease notation, in what follows we drop indices on $x_i(j)$ and represent an ISI generically by x.

One term which occurs several times in defining the accumulated evidence is $\log x$. Now, if r is the instantaneous firing rate with $r = 1/x$, then $\log x = -\log r$. The transfer function, from current z to firing rate w, of the simple leaky integrate and fire neuron (LIF) can, under the correct circumstances, approximate the form $w = log(z)$ [88]. Now suppose that the input spikes with rate r give, after low-pass filtering by the membrane, a roughly constant current z. This will be proportional to r, as $z = cr$, and so we have $w = log(cr) = c' + log(r)$, with $c' = log(c)$. The constant terms like c' may be absorbed into the term $B(T)$ in Eq 8. Thus if r was an input to an inhibitory LIF-like neuron, the neuron's output could form an additive input $-log(r)$ to a second neural stage which combines the required terms.

Now consider the computation of terms like $(\log x)^2$ which occur with the use of the lognormal distribution. Recall from the narrative immediately following Eq 11, that the sum of ISIs to decision time for the i^{th} channel is just T_i, and that this is approximately T for all i. Further, such constant terms makes no overall contribution to the outcome. Thus, a term linear in x, which gives rise to such a sum of ISIs, will have no effect on the decision time (within the approximation $T = T_i$). We can, therefore, always add a term linear in x to any component with negligible effect on the decision outcome. Thus, we can replace the squared log term by $(\log x)^2 - a^2 x$ with constant a. Using $x = 1/r$, this becomes $(-\log r + a/\sqrt{r})(-\log r - a/\sqrt{r})$. Each bracket contains a term in $-\log r$ which we can compute using the procedure described above. The term $a/\sqrt{r}$ involves division by $\sqrt{r}$ and could be achieved by shunting inhibition acting on a tonic firing rate a [89–91]. The multiplication of the brackets is also plausible using active dendritic processing [92, 93].

Finally, terms linear in x can (as noted previously) be ignored, and terms like $1/x$ correspond to additive input based on firing rate—often the default assumption for neural processing. In sum, all the computations required for $L_i(j)$ in Table 1 are available to neural machinery at the synaptic and circuit levels.

From ISIs to behaviour

The s-MSPRT provides a direct link from the information in spike train ISIs to 'neural decision times' which are a component of an overall behavioural decision time or reaction time. Further this link is strengthened by the hypothesis, supported in our simulation work, that the mean total information gain required to make a decision is constant (for a given error rate and number of alternatives). This was a key to providing explanations for Hick's and Piéron's laws and provides a direct link from neural signalling to psychophysical observation. We argue that we have laid the foundations for a programme of work in which neural recordings and spike trains' statistical analysis can be made the basis of predictions about behaviour in multi-alternative choice tasks.

Methods

Choosing probability distributions to model spike data

Fig 9 shows the distribution of ISIs (grey histogram) in a macaque MT neuron, from the study by Britten et al [30, 31] using the RDMT. In this task, the monkey is typically shown two eccentric targets and then a kinematogram composed of dots in which a given percentage of them (the so-called 'coherence') move towards one of the targets, while the others move randomly. The animal is rewarded if it saccades towards this target from a gaze fixation point [30, 94]. The RDMT is therefore a perceptual decision making task of the kind we envisage being solved by a mechanism like the one proposed (s-MSPRT).

Each panel of Fig 9 shows the best fit (using the method of moments) of a pdf to the experimental data in the histogram. Data is assumed to come from single neurons and thus to be drawn from one such pdf; the implications for multiple neurons are considered further in the Discussion. All the pdfs are from the exponential family. The choice of putative distributions had the following rationale. The Gaussian, while an oft-used choice in more abstract (*e.g.* machine learning) approaches, cannot, in principle, be a genuine candidate for ISIs as there will always be a negative tail to the distribution which is physically implausible (there is no such thing as a 'negative ISI'). However, it might be argued that this tail is, in most realistic cases, very small and that it can be neglected. We tested this possibility with real data sets (including that shown in Fig 9). The exponential distribution of ISIs leads to spike trains following a Poisson process—another popular choice for spike train analysis. This time, while the distribution satisfies the positivity requirement, it does not satisfy the requirement for excluding arbitrarily small ISIs due to a neural refractory period. In order to achieve this the distribution must have positive mode as well as lying wholly in the positive half-plane. In addition, we require that the pdf is, in general, positively skewed which matches the general shape observed in ISI data. These requirements are satisfied by suitable parameterisations of the other four candidate distributions: lognormal, gamma, inverse-Gaussian and inverse gamma.

The inverse Gaussian and lognormal appear to be the best fit to the single data set in Fig 9. In contrast, note the large (implied) negative tail of the Gaussian, which could certainly not be neglected in using it to sample from for a decision process. However, aside from these qualitative observations, we sought to quantify the fitting process by repeating it with several data sets, and using two 'goodness of fit' metrics, the Kolmogorov-Smirnov and Anderson-Darling (see Fig 10) [95]. These tests are among the most powerful available for continuous and completely specified distributions [96]. The datasets are taken mainly from the Britten et al study [30] and are the first five described in Table 2. In order to avoid the transient due to stimulus onset, as advised by Bair [21] working with the same data set, we discarded the first 336 ms of all used spike trains. There is one other data set comprising recordings in the MT/middle-superior-
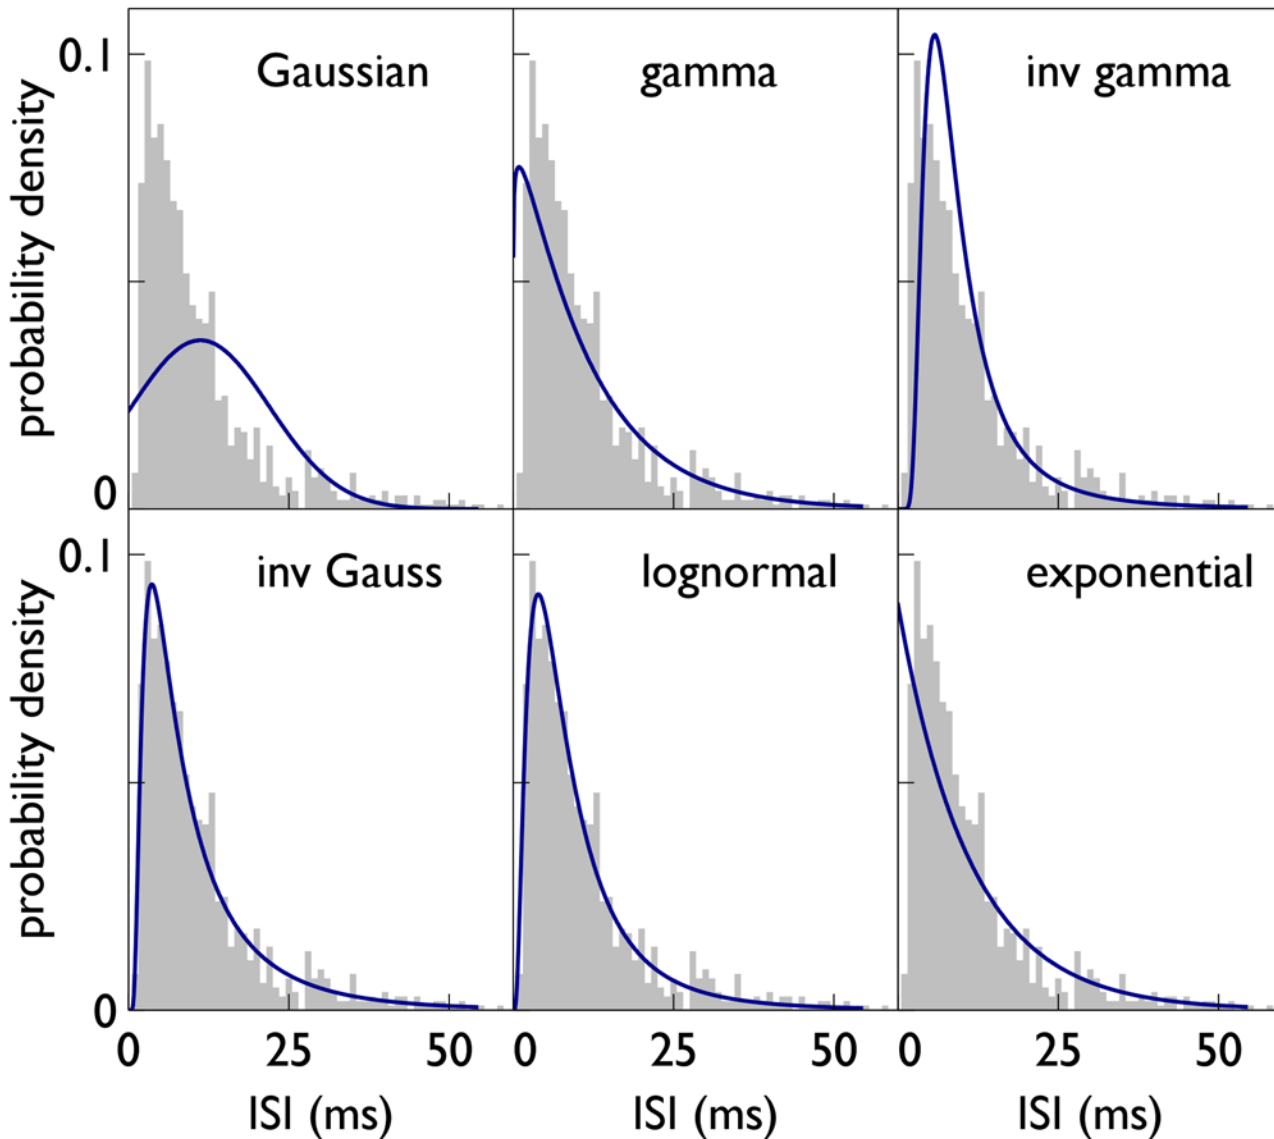
Fig 9. Fitting pdfs to an ISI histogram. The ISIs in the grey histogram (identical in each panel) were recorded by [31] from the MT neuron with tag e093 (Table 2). The few ISIs lying farther than four standard deviations beyond the mean were not plotted for clarity. Overlaid as a solid blue line in each panel is the best-fit pdf from a set of theoretical distributions: Gaussian, gamma, inverse gamma, inverse Gaussian, lognormal and exponential

doi:10.1371/journal.pone.0124787.g009

temporal area, made during a paradigm using visual motion cues to initiate arm movements [27].

Fitting results are in Fig 10, the smaller the statistic the better the fit. It is clear that the Gaussian is the worst fit to these data. Indeed, the scale on the Anderson-Darling plot is logarithmic largely to accommodate the Gaussian fits which have a mean of 62.5 compared to the next worst fit (the gamma) whose mean is 20.4. The exponential and gamma distribution had intermediate levels of fit, but were chosen for further study because of the useful comparison to be made with the oft-used exponential (giving Poisson statistics), and the fact that the gamma simplifies to the exponential when its mean and standard deviation are equal. The last three distributions (inverse gamma, lognormal and inverse-Gaussian) all have similar mean
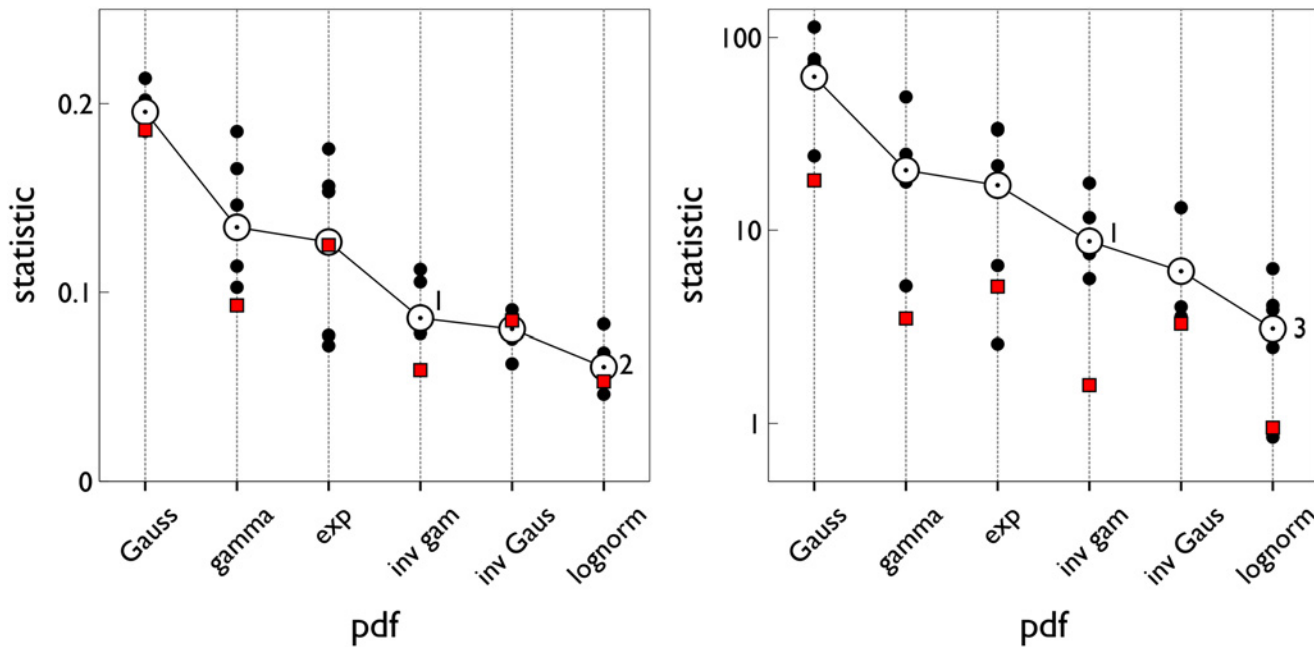
Fig 10. Goodness of fit of selected pdfs to ISI data. The black closed circles show the goodness of fit statistics for each of the first five data sets in Table 2. The red closed squares are for the data set from [27]. The mean for all six data sets is shown by the large open circles which also has a line plot; the number of significant fits at a level of 0.05 is noted per pdf next to such circles. Panels a and b are for the Kolomogorov-Smirnov, and Anderson-Darling tests respectively (note log y-axis in the latter). The ordering of the results from left to right preserves rank order of the mean statistic and is the same for both tests.

doi:10.1371/journal.pone.0124787.g010

statistics, and are the best fits of all; in fact, it is remarkable that in a few instances (numbers next to the average-statistic symbols) the inverse gamma and lognormal pdfs significantly fitted the experimental data. In sum, we selected the gamma, inverse gamma, lognormal, inverse-Gaussian and exponential distributions for further study.

Parameterising the distributions

For each selection of distribution we have to set four parameters: the mean μ^* and standard deviation σ^* of the data and test distributions f_* of the preferred channel or hypothesis (respectively), and their null counterparts μ_0, σ_0 for f_0. We used two base-parameter sets, from which others were derived, The first base set we dub the *independent variance* parameter set, Ω_{IV}, as both means and variances were independently fitted to one of the data sets from the RDMT study by Britten et al [30, 31]. Thus, μ^*, σ^* were fitted to the responses of neuron w144 when this was stimulated at coherence 12.8% with dots moving in its preferred direction (last row in Table 2). This neuron was singled out for being 'typical' of its parent population as its statistics closely matched a mean of the ISIs of over 40 MT neurons from [31]. The value of μ_0, σ_0 were obtained by fitting to the same neuron when stimulated at the same coherence but with motion in the null (anti-preferred) direction (sixth row in Table 2). This gave

$$\Omega_{IV} = \quad \{\mu^* = 16.5, \quad \sigma^* = 21.5, \quad \mu_0 = 33, \quad \sigma_0 = 47.5\}, \quad \text{all units ms} \quad (32)$$

The distributions resulting from using Ω_{IV} are shown in Fig 3a–3d. Other parameter sets (such as those in Fig 5a) are derived from Ω_{IV} by changing μ_0, but keeping $\mu_0 - \mu^*$ constant, and keeping the standard deviations of Ω_{IV}. We will refer to these sets as $\hat{\Omega}_{IV}(\mu_0)$

Table 2. Data sets from [31] used to generate ISI distributions.

Neuron tag	coherence	direction	number of trials	number of ISIs
j200	99.9%	1	7	1138
j200	6.4%	0	16	1030
e093*	25.6%	0	7	1044
e093	12.8%	1	6	1042
w059	51.2%	1	7	341
w144	12.8%	0	30	1360
w144	12.8%	1	30	2790

The Neuron tag is that used in the original data repository; the asterisk denotes the data used for the histogram in Fig 9. 'Coherence' is the percentage of the dots moving coherently in the random dot stimulus. 'Direction' (0 or 1) indicates overall stimulus direction: direction '1' is the preferred direction of the neuron (across its entire tuning curve) while direction '0' is its opposite. 'Number of ISIs' is the size of our sample, resulting from pooling data across the 'number of trials' indicated in the adjacent column.

doi:10.1371/journal.pone.0124787.t002

The exponential cannot be described by a set like Ω_{IV} as it only has a single parameter $\lambda = 1/\mu$. In order to make comparisons with the exponential distribution, we therefore defined another *fixed variance* base-parameter set Ω_{FV}

$$\Omega_{FV} = \{\mu^* = 16.5, \quad \sigma^* = 16.5, \quad \mu_0 = 33, \quad \sigma_0 = 33\}, \quad \text{all units ms} \tag{33}$$

Other parameter sets (such as those in Fig 5b) are obtained by changing μ_0, but keeping $\mu_0 - \mu^*$ constant, and varying the standard deviations appropriately in order to maintain compatibility with the exponential. We refer to these sets as $\hat{\Omega}_{FV}(\mu_0)$.

Bounds on mean total information gain

Here, we give a plausibility argument for the empirical result in Eq 22. It is based on Lemma 2 in [44] which is framed in a more general situation where each hypothesis corresponds to a probability distribution f_j. The Lemma makes uses the error probabilities

$$\alpha_{k,l} = P(\text{decide } H_l \text{ is true}|H_k \text{ is true}) \tag{34}$$

and states that, the expected number of observations $\langle s \rangle_k$ to decision that any hypothesis k is true (with respect to a reference j), is given by

$$\langle s \rangle_k \geq \frac{1}{D(f_k||f_j)} \sum_{l=1}^{N} \alpha_{k,l} \log \frac{\alpha_{k,l}}{\alpha_{j,l}} \tag{35}$$

In our situation, if ϵ is the error rate of the decision process,

$$\alpha_{j,j} = 1 - \epsilon, \qquad \alpha_{j,l} = \frac{\epsilon}{N-1}, \quad j \neq l \tag{36}$$

so the sum on the right hand side of Eq 35 has only two non-zero terms

$$
\begin{aligned}
\sum_{l=1}^{N} \alpha_{k,l} \log \frac{\alpha_{k,l}}{\alpha_{j,l}} &= (1-\epsilon) \log \frac{(1-\epsilon)(N-1)}{\epsilon} + \frac{\epsilon}{N-1} \log \frac{\epsilon}{(1-\epsilon)(N-1)} \\
&= \left(1 - \epsilon - \frac{\epsilon}{N-1}\right) \log \frac{(1-\epsilon)(N-1)}{\epsilon} \\
&\equiv \mathcal{E}(\epsilon, N)
\end{aligned}
\tag{37}
$$

Then, using our notation for the two distributions and for the expected number of observations to decision, Eq 35 for the preferred channel becomes

$$
\langle s_s^* \rangle \geq \frac{\mathcal{E}(\epsilon, N)}{D(f_* \| f_0)}
\tag{38}
$$

We now compare this result with that in Eq 20. Being an inequality, it might appear at first that Eq 38 is a weaker result than the equality in Eq 20. However, Eq 38 it is a strong result in so far as the numerator is explicitly independent of the distributions or their parameterisation. This makes it plausible that, in Eq 20, $\langle y_*(T_D^s) \rangle_{f_*}$ is relatively constant. We therefore identify $\mathcal{E}$ with A in Eq 22, so that A is now a function of ϵ and N with functional form given in Eq 37. Further, using the values of N and ϵ in our simulations, $\mathcal{E}(0.05, 10) = 4.856$ which is only 5.7% different from the empirical value of 4.594 for the constant A in Eq 22. Taking these observations together, we conclude that, at least over the values of N and ϵ comparable with those used here, Eq 22 and Eq 24 have a plausible basis.

A heuristic analysis of the relation between s-MSPRT and u-MSPRT

Returning to the analysis of $R(T)$, let $\langle y_0^u(T) \rangle$, $\langle y_0^s(T) \rangle$ be the expected value of the evidence on a null channel for u- and s-MSPRT, respectively. Then, using the arguments above, $\langle y_0^u(T) \rangle$ is formed from more observations than $\langle y_0^s(T) \rangle$, and their contributions are likely to be negative. Thus, $\langle y_0^u(T) \rangle < \langle y_0^s(T) \rangle$. Taking this result together with observation matching on the preferred hypothesis it is plausible that

$$
\langle \exp y_0^u(T) \rangle < \langle \exp y_0^s(T) \rangle, \quad \langle \exp y_*^u(T) \rangle = \langle \exp y_*^s(T) \rangle
\tag{39}
$$

Let $\langle R_u(T) \rangle$, $\langle R_s(T) \rangle$ be the expected values of the normalisation terms in u- and s-MSPRT, where

$$
\langle R_u(T) \rangle = \langle \exp y_*^u(T) \rangle + (N-1)\langle \exp y_0^u(T) \rangle, \quad \langle R_s(T) \rangle = \langle \exp y_*^s(T) \rangle + (N-1)\langle \exp y_0^s(T) \rangle
\tag{40}
$$

Then using Eq 39, $\langle R_u(T) \rangle < \langle R_s(T) \rangle$. The expected value of the posterior $P_*(T)$ for u-MSPRT will therefore be larger at T than that for s-MSPRT. We conclude that, in order to preserve the observation count in the preferred channel before crossing threshold, we would expect the threshold θ_u for u-MSPRT to be greater than that for s-MSPRT, θ_s. This result carries over to any monotonic function of the thresholds such as the exponential.

Fig 11 shows that this condition on the thresholds is indeed met. Fig 11a and 11b show exp$(\theta_s) < \exp(\theta_u)$ consistently across all numbers of choices for two examples (lognormal and inverse Gaussian based s-MSPRT) for parameter set Ω_{IV}. (The lognormal is more typical with little variation with the number of alternatives). Similar clear differences $\Delta_e \theta = \exp(\theta_u) - \exp(\theta_s)$ exist for all distributions across both parameter sets Ω_{IV}, Ω_{FV}—see Fig 11c and 11d. Notice that the gamma distribution in Fig 11c and the exponential distribution in Fig 11d have the smallest $\Delta_e \theta$ in their groups (and in each case, $\Delta_e \theta$ is significantly less than the next smallest one at $p < 0.001$, two-sided t-test). This may account for the fact that decision times for the gamma
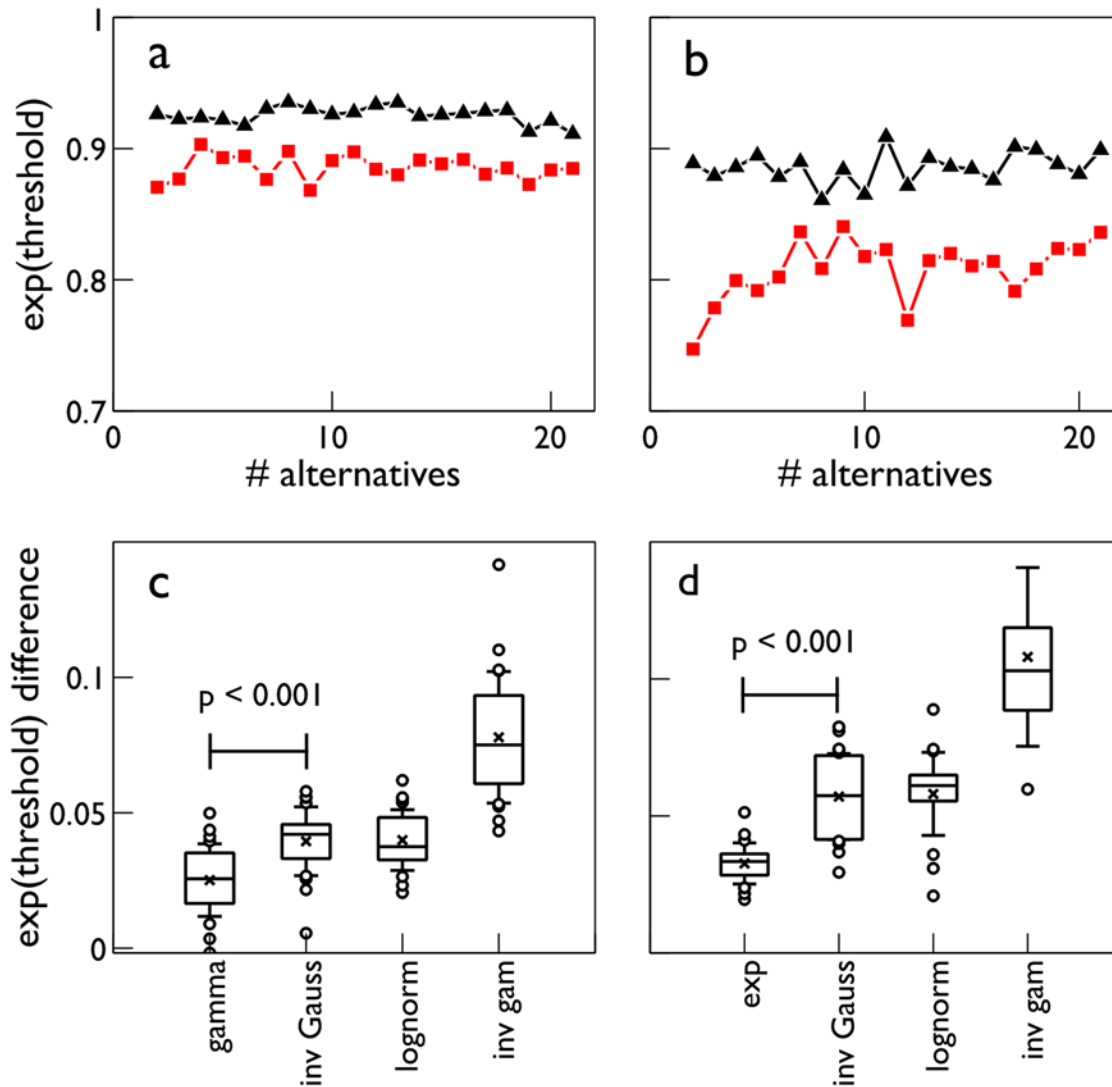
Fig 11. Decision thresholds and their difference between s-MSPRT and u-MSPRT. To ease visualisation, all panels show the exponential of the thresholds; $\phi_s = \exp(\theta_s)$ $\phi_u = \exp(\theta_u)$ for s- and u-MSPRT respectively. This yields positive values pertaining to the posterior (rather than negative values for the log-posterior). Panels a and b, respectively, show the lognormal and inverse Gaussian cases in Fig 4a for the parameter set Ω_{IV}. The red lines and symbols are for ϕ_s, the black lines and symbols for ϕ_u. In panel c, the box plot labelled 'lognormal' shows the median and quartiles (box lines), mean (cross) and one standard deviation of the differences $\phi_u - \phi_s$ in panels a and b. Other bars show similar quantities for the test distributions of the other MSPRT instantiations used to form Fig 4a. Panel d is similar to panel c, except it pertains to differences in (exponential) thresholds for results in Fig 4b, with the parameter set Ω_{FV}.

doi:10.1371/journal.pone.0124787.g011

distribution (Fig 4a) and the exponential distribution (Fig 4b) fail to achieve the criterion of equal decision sample as well as the others.

Acknowledgments

We would like to thank Mark D. Humphries, Rafal Bogacz, Jim V. Stone and Andreas Steimer for their invaluable discussion on some of the results of this work.

Author Contributions

Conceived and designed the experiments: JC NL KG. Performed the experiments: JC KG. Analyzed the data: JC NL KG. Contributed reagents/materials/analysis tools: JC KG. Wrote the paper: JC NL KG.

References

1. Ratcliff R (1978) A theory of memory retrieval. Psychological Review 85: 59–108. doi: 10.1037/0033-295X.85.2.59

2. Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. Nature 377: 59–62. doi: 10.1038/377059a0 PMID: 7659161

3. Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. Cereb Cortex 13: 1257–1269. doi: 10.1093/cercor/bhg097 PMID: 14576217

4. Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev Neurosci 30: 535–574. doi: 10.1146/annurev.neuro.29.051605.113038 PMID: 17600525

5. Rao RPN (2010) Decision making under uncertainty: a neural model based on partially observable markov decision processes. Front Comput Neurosci 4: 146. doi: 10.3389/fncom.2010.00146 PMID: 21152255

6. Churchland AK, Ditterich J (2012) New advances in understanding decisions among multiple alternatives. Current opinion in neurobiology 22: 920–926. doi: 10.1016/j.conb.2012.04.009 PMID: 22554881

7. Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. Trends Cogn Sci 5: 10–16. doi: 10.1016/S1364-6613(00)01567-9 PMID: 11164731

8. McMillen T, Holmes P (2006) The dynamics of choice among multiple alternatives. Journal of Mathematical Psychology 50: 30–57. doi: 10.1016/j.jmp.2005.10.003

9. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol Rev 113: 700–765. doi: 10.1037/0033-295X.113.4.700 PMID: 17014301

10. Bogacz R, Gurney K (2007) The basal ganglia and cortex implement optimal decision making between alternative actions. Neural Comput 19: 442–477. doi: 10.1162/neco.2007.19.2.442 PMID: 17206871

11. Gerstein GL, Mandelbrot B (1964) Random walk models for the spike activity of a single neuron. Biophysical Journal 4: 41–68. doi: 10.1016/S0006-3495(64)86768-0 PMID: 14104072

12. Smith PL (2010) From poisson shot noise to the integrated Ornstein-Uhlenbeck process: Neurally principled models of information accumulation in decision-making and response time. Journal of Mathematical Psychology 54: 266–283. doi: 10.1016/j.jmp.2009.12.002

13. Purcell BA, Heitz RP, Cohen JY, Schall JD, Logan GD, et al. (2010) Neurally constrained modeling of perceptual decision making. Psychological review 117: 1113. doi: 10.1037/a0020311 PMID: 20822291

14. Purcell BA, Schall JD, Logan GD, Palmeri TJ (2012) From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. The Journal of Neuroscience 32: 3433–3446. doi: 10.1523/JNEUROSCI.4622-11.2012 PMID: 22399766

15. Beck JM, Ma W, Kiani R, Hanks T, Churchland AK, et al. (2008) Probabilistic population codes for bayesian decision making. Neuron 60: 1142–1152. doi: 10.1016/j.neuron.2008.09.021 PMID: 19109917

16. Ditterich J (2010) A comparison between mechanisms of multi-alternative perceptual decision making: Ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. Front Neurosci 4: 184. doi: 10.3389/fnins.2010.00184 PMID: 21152262

17. Zhang J, Bogacz R (2010) Optimal decision making on the basis of evidence represented in spike trains. Neural Comput 22: 1113–1148. doi: 10.1162/neco.2009.05-09-1025 PMID: 20028228

18. Steimer A, Douglas R (2013) Spike-based probabilistic inference in analog graphical models using interspike-interval coding. Neural computation 25: 2303–2354. doi: 10.1162/NECO_a_00477 PMID: 23663144

19. Lepora NF, Gurney KN (2012) The basal ganglia optimize decision making over general perceptual hypotheses. Neural Comput 24: 2924–2945. doi: 10.1162/NECO_a_00360 PMID: 22920846

20. Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. J Neurosci 13: 334–350. PMID: 8423479

21. Bair W, Koch C, Newsome W, Britten K (1994) Power spectrum analysis of bursting cells in area MT in the behaving monkey. J Neurosci 14: 2870–2892. PMID: 8182445

22. Baddeley R, Abbott LF, Booth MC, Sengpiel F, Freeman T, et al. (1997) Responses of neurons in primary and inferior temporal visual cortices to natural scenes. Proceedings of the Royal Society of London Series B: Biological Sciences 264: 1775–1783. doi: 10.1098/rspb.1997.0246 PMID: 9447735

23. Barbieri R, Quirk MC, Frank LM, Wilson MA, Brown EN (2001) Construction and analysis of non-Poisson stimulus-response models of neural spiking activity. J Neurosci Methods 105: 25–37. doi: 10.1016/S0165-0270(00)00344-7 PMID: 11166363

24. de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. Nat Neurosci 8: 1698–1703. doi: 10.1038/nn1587 PMID: 16286929

25. Churchland MM, Byron MY, Ryu SI, Santhanam G, Shenoy KV (2006) Neural variability in premotor cortex provides a signature of motor preparation. The Journal of neuroscience 26: 3697–3712. doi: 10.1523/JNEUROSCI.3762-05.2006 PMID: 16597724

26. Hromádka T, DeWeese MR, Zador AM (2008) Sparse representation of sounds in the unanesthetized auditory cortex. PLoS Biol 6: e16. doi: 10.1371/journal.pbio.0060016 PMID: 18232737

27. Maimon G, Assad JA (2009) Beyond Poisson: increased spike-time regularity across primate parietal cortex. Neuron 62: 426–440. doi: 10.1016/j.neuron.2009.03.021 PMID: 19447097

28. Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. Trends Neurosci 27: 161–168. doi: 10.1016/j.tins.2004.01.006 PMID: 15036882

29. Baum C, Veeravalli V (1994) A sequential procedure for multihypothesis testing. IEEE Transactions on Information Theory 40: 1994–2007. doi: 10.1109/18.340472

30. Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci 12: 4745–4765. PMID: 1464765

31. Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992). Responses of single neurons in macaque MT/V5 as a function of motion coherence in stochastic dot stimuli, Neur Sig Arch, NSA2004.1. www.neuralsignal.org. URL http://www.neuralsignal.org/.

32. Beck JM, Ma WJ, Pitkow X, Latham PE, Pouget A (2012) Not noisy, just wrong: the role of suboptimal inference in behavioral variability. Neuron 74: 3039. doi: 10.1016/j.neuron.2012.03.016

33. Kullback S, Leibler RA (1951) On information and sufficiency. The Annals of Mathematical Statistics 22: 79–86. doi: 10.1214/aoms/1177729694

34. Kullback S (1987) Letter to the editor: The Kullback-Leibler distance. The American Statistician 41: 340–341.

35. Hick WE (1952) On the rate of gain of information. The Quarterly Journal of Experimental Psychology 4: 11–26. doi: 10.1080/17470215208416600

36. Pieron H (1913) II. recherches sur les lois de variation des temps de latence sensorielle en fonction des intensites excitatrices. L'annee psychologique 20: 17–96. doi: 10.3406/psy.1913.4294

37. Caballero JA, Gurney KN (2010) Optimal decision making with biologically realistic neural signals. BMC Neuroscience 11: P25. doi: 10.1186/1471-2202-11-S1-P25

38. Caballero JA, Lepora N, Gurney KN (2013) Decision-making out of neural events: from discrimination information to psychometric power laws. BMC Neuroscience 14: P153. doi: 10.1186/1471-2202-14-S1-P153

39. Caballero JA (2012) Multi-alternative sequential analysis as a realistic model of biological decision-making. PhD thesis, The University of Sheffield, Sheffield, UK.

40. Dragalin VP, Tartakovsky AG, Veeravalli VV (2000) Multihypothesis sequential probability ratio tests. II. accurate asymptotic expansions for the expected sample size. IEEE Trans Inform Theory 46: 1366–1383. doi: 10.1109/18.850677

41. Press WH, Teukolsky SA, Vetterling WT, Flannery BP (2007) Numerical recipes: The art of scientific computing. Cambridge university press, 3rd edition.

42. Wald A (1944) On cumulative sums of random variables. The Annals of Mathematical Statistics 15: 283296. doi: 10.1214/aoms/1177731235

43. Bruss FT, Robertson JB (1991) 'Wald's lemma' for sums of order statistics of iid random variables. Advances in applied probability: 612–623. doi: 10.2307/1427625

44. Veeravalli W, Baum CW (1995) Asymptotic efficiency of a sequential multihypothesis test. IEEE Transactions on Information Theory 41: 1994–1997. doi: 10.1109/18.476323

45. Hyman R (1953) Stimulus information as a determinant of reaction time. Ph.D. thesis, Johns Hopkins University. URL http://www.psychology.uiowa.edu/faculty/mordkoff/InfoProc/pdfs/Hyman%201953.pdf.

46. Seow SC (2005) Information theoretic models of HCI: a comparison of the Hick-Hyman law and Fitts' law. Human-Computer Interaction 20: 315352. doi: 10.1207/s15327051hci2003_3

47. Usher M, Olami Z, McClelland JL (2002) Hick's law in a stochastic race model with speedaccuracy tradeoff. Journal of Mathematical Psychology 46: 704715. doi: 10.1006/jmps.2002.1420

48. Leonard JA (1959) Tactual choice reactions: I. Quarterly Journal of Experimental Psychology 11: 76–83. doi: 10.1080/17470215908416294

49. Kveraga K, Boucher L, Hughes HC (2002) Saccades operate in violation of Hicks law. Experimental Brain Research 146: 307–314. doi: 10.1007/s00221-002-1168-8 PMID: 12232687

50. Piéron H (1952) The Sensations: their functions, processes and mechanisms. London, UK: Frederick Muller, Ltd.

51. Pins D, Bonnet C (1996) On the relation between stimulus intensity and processing time: Piéron's law and choice reaction time. Perception and Psychophysics 58: 390–400. doi: 10.3758/BF03206815 PMID: 8935900

52. Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. Journal of vision 5: 1. doi: 10.1167/5.5.1

53. Stafford T, Ingram L, Gurney KN (2011) Péron's law holds during stroop conflict: Insights into the architecture of decision making. Cognitive science 35: 15531566. doi: 10.1111/j.1551-6709.2011.01195.x

54. Stafford T, Gurney KN (2004) The role of response mechanisms in determining reaction time performance: Pirons law revisited. Psychonomic bulletin & review 11: 975987.

55. van Maanen L, Grasman RPPP, Forstmann BU, Wagenmakers EJ (2012) Piéron's law and optimal behavior in perceptual decision-making. Front Neurosci 5: 143. doi: 10.3389/fnins.2011.00143 PMID: 22232572

56. Muniak MA, Ray S, Hsiao SS, Dammann JF, Bensmaia SJ (2007) The neural coding of stimulus intensity: linking the population response of mechanoreceptive afferents with psychophysical behavior. The Journal of Neuroscience 27: 11687–11699. doi: 10.1523/JNEUROSCI.1486-07.2007 PMID: 17959811

57. Vickers D (1970) Evidence for an accumulator model of psychophysical discrimination. Ergonomics 13: 37–58. doi: 10.1080/00140137008931117 PMID: 5416868

58. Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. Psychol Rev 108: 550–592. doi: 10.1037/0033-295X.108.3.550 PMID: 11488378

59. Cisek P, Kalaska JF (2010) Neural mechanisms for interacting with a world full of action choices. Annu Rev Neurosci 33: 269–298. doi: 10.1146/annurev.neuro.051508.135409 PMID: 20345247

60. Resulaj A, Kiani R, Wolpert DM, Shadlen MN (2009) Changes of mind in decision-making. Nature 461: 263–266. doi: 10.1038/nature08275 PMID: 19693010

61. Schall JD, Hanes DP (1993) Neural basis of saccade target selection in frontal eye field during visual search. Nature 366: 467–469. doi: 10.1038/366467a0 PMID: 8247155

62. Horwitz GD, Newsome WT (1999) Separate signals for target selection and movement specification in the superior colliculus. Science 284: 1158–1161. doi: 10.1126/science.284.5417.1158 PMID: 10325224

63. Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat Neurosci 2: 176–185. doi: 10.1038/5739 PMID: 10195203

64. Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci 22: 9475–9489. PMID: 12417672

65. Hernández A, Nácher V, Luna R, Zainos A, Lemus L, et al. (2010) Decoding a perceptual decision process across cortex. Neuron 66: 300–314. doi: 10.1016/j.neuron.2010.03.031 PMID: 20435005

66. Ding L, Gold JI (2010) Caudate encodes multiple computations for perceptual decisions. J Neurosci 30: 15747–15759. doi: 10.1523/JNEUROSCI.2894-10.2010 PMID: 21106814

67. Ding L, Gold JI (2012) Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. Cerebral Cortex 22: 1052–1067. doi: 10.1093/cercor/bhr178 PMID: 21765183

68. Stanford TR, Shankar S, Massoglia DP, Costello MG, Salinas E (2010) Perceptual decision making in less than 30 milliseconds. Nat Neurosci 13: 379–385. doi: 10.1038/nn.2485 PMID: 20098418

69. Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. Nat Neurosci 11: 693–702. doi: 10.1038/nn.2123 PMID: 18488024

70. Niwa M, Ditterich J (2008) Perceptual decisions between multiple directions of visual motion. The Journal of Neuroscience 28: 4435–4445. doi: 10.1523/JNEUROSCI.5564-07.2008 PMID: 18434522

71. Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A (2012) The cost of accumulating evidence in perceptual decision making. J Neurosci 32: 3612–3628. doi: 10.1523/JNEUROSCI.4010-11.2012 PMID: 22423085

72. Lindner B (2006) Superposition of many independent spike trains is generally not a poisson process. Physical Review E 73: 022901. doi: 10.1103/PhysRevE.73.022901

73. Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. Annual review of neuroscience 21: 227277. doi: 10.1146/annurev.neuro.21.1.227

74. Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, et al. (2014) Information-limiting correlations. Nature neuroscience. doi: 10.1038/nn.3807 PMID: 25195105

75. Nawrot MP, Boucsein C, Rodriguez Molina V, Riehle A, Aertsen A, et al. (2008) Measurement of variability dynamics in cortical spike trains. Journal of neuroscience methods 169: 374–390. doi: 10.1016/j.jneumeth.2007.10.013 PMID: 18155774

76. Chacron MJ, Lindner B, Longtin A (2004) Noise shaping by interval correlations increases information transfer. Physical review letters 92: 080601. doi: 10.1103/PhysRevLett.92.080601 PMID: 14995762

77. Wald A (1947) Sequential analysis. Wiley publications in mathematical statistics. New York: John Wiley.

78. Recanzone GH, Merzenich MM, Jenkins WM, Grajski KA, Dinse HR (1992) Topographic reorganization of the hand representation in cortical area 3b owl monkeys trained in a frequency-discrimination task. J Neurophysiol 67: 1031–1056. PMID: 1597696

79. Recanzone Ga, Schreiner C, Merzenich MM (1993) Plasticity in the frequency representation of primary auditory cortex following discrimination training in adult owl monkeys. The Journal of Neuroscience 13: 87–103. PMID: 8423485

80. Bao S, Chang EF, Woods J, Merzenich MM (2004) Temporal plasticity in the primary auditory cortex induced by operant perceptual learning. Nature neuroscience 7: 974–981. doi: 10.1038/nn1293 PMID: 15286790

81. Jeyabalaratnam J, Bharmauria V, Bachatene L, Cattan S, Angers A, et al. (2013) Adaptation shifts preferred orientation of tuning curve in the mouse visual cortex. PloS one 8: e64294. doi: 10.1371/journal.pone.0064294 PMID: 23717586

82. Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nature neuroscience 11: 505–513. doi: 10.1038/nn2070 PMID: 18327253

83. Laughlin S (1981) A simple coding procedure enhances a neuron's information capacity. Z Naturforsch C 36: 910–912. PMID: 7303823

84. Fairhall AL, Lewen GD, Bialek W, van Steveninck RRdR (2001) Efficiency and ambiguity in an adaptive neural code. Nature 412: 787–792. doi: 10.1038/35090500 PMID: 11518957

85. Nemenman I, Bialek W, van Steveninck RdR (2004) Entropy and information in neural spike trains: Progress on the sampling problem. Physical Review E 69: 056111. doi: 10.1103/PhysRevE.69.056111

86. Bogacz R, Larsen T (2011) Integration of reinforcement learning and optimal decision-making theories of the basal ganglia. Neural Comput. doi: 10.1162/NECO_a_00103

87. Caballero JA, Lepora N, Gurney KN (2011) Sequential tests and biologically grounded multi-alternative decision making. BMC Neuroscience 12: P137. doi: 10.1186/1471-2202-12-S1-P137

88. Tal D, Schwartz EL (1997) Computing with the leaky integrate-and-fire neuron: logarithmic computation and multiplication. Neural Computation 9: 305–318. doi: 10.1162/neco.1997.9.2.305 PMID: 9117905

89. Prescott SA, De Konink Y (2003) Gain control of firing rate by shunting inhibition: Roles of synaptic noise and dendritic saturation. Proc Natl Acad Sci USA 100: 2076–81. doi: 10.1073/pnas.0337591100 PMID: 12569169

90. Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. Neuron 35: 773–82. doi: 10.1016/S0896-6273(02)00820-6 PMID: 12194875

91. Mehaffey WH, Doiron B, Maler L, Turner RW (2005) Deterministic multiplicative gain control with active dendrites. Journal of Neuroscience 25: 9968–77. doi: 10.1523/JNEUROSCI.2682-05.2005 PMID: 16251445

92.   London M, Häusser M (2005) Dendritic computation. Annu Rev Neurosci 28: 503–532. doi: 10.1146/annurev.neuro.28.061604.135703 PMID: 16033324

93.   Mel B, Ruderman D, Archie K (1998) Translation-invariant orientation tuning in visual "complex" cells could derive from intradendritic computations. J Neuroscience 18: 4325–4334.

94.   Shadlen MN, Newsome WT (1996) Motion perception: seeing and deciding. Proc Natl Acad Sci U S A 93: 628–633. doi: 10.1073/pnas.93.2.628 PMID: 8570606

95.   Thas O (2009) Comparing distributions. Springer Series in Statistics. Springer.

96.   Stephens MA (1974) EDF statistics for goodness of fit and some comparisons. Journal of the American Statistical Association 69: 730–737. doi: 10.1080/01621459.1974.10480196