Achieving Constant Hazard Rate in Psychophysics Experiments

Authors:

Brian Maniscalco[1] & Megan A. K. Peters[1,2,3,4]

Affiliations:

[1] Department of Cognitive Sciences, University of California Irvine, Irvine, CA

[2] Center for the Neurobiology of Learning & Memory, University of California Irvine, Irvine, CA

[3] Center for Theoretical Behavioral Sciences, University of California Irvine, Irvine, CA

[4] Program in Brain, Mind, & Consciousness, Canadian Institute of Advanced Research, Toronto, Ontario

Correspondence should be directed to:

Brian Maniscalco (bmaniscalco@gmail.com)

Megan A. K. Peters (megan.peters@uci.edu)

Abstract

In psychophysics, it is sometimes desirable for the temporal onset of an event to be unpredictable, e.g. to control possible effects of expectation and temporal attention in behavioral or neural data. One way of achieving temporal unpredictability is to make the *hazard rate* of the event constant. Roughly speaking, the hazard rate measures the odds that an event will occur now or later. If the hazard rate is constant throughout a temporal interval, then the time elapsed thus far provides the observer with no information for updating their expectation of when the event will occur. A constant hazard rate thus prevents the building of predictive anticipation that the event is about to occur even as time continues to unfold. Here we consider the question of how best to control hazard rates in psychophysics experiments with time considered as either a continuous or discrete variable. We argue that for computerized psychophysics experiments, it is both more appropriate and more convenient to model time as a discrete variable. We also consider that for practical reasons, experimenters may wish to impose upper bounds on the duration of random temporal intervals. We show how hazard rate can be exactly controlled for unbounded time intervals when time is treated as either a continuous or discrete variable by using the exponential and geometric distributions, respectively. We discuss complications in controlling hazard rate for bounded time intervals and present a solution yielding exact control of hazard rate for bounded time intervals with discrete time steps using a modified form of the geometric distribution. The latter approach is likely to be the most relevant for most psychophysics experiments that seek to control hazard rate.

Introduction

In the design of psychophysics experiments, it is sometimes desirable for the temporal onset of certain events that occur repeatedly over many trials to be unpredictable, e.g. to control possible effects of expectation and temporal attention in behavioral or neural data. One might initially suppose that randomly drawing the time until event occurrence on each trial from a uniform distribution over $t \in [0, b]$ for some upper time limit b would be a natural way to achieve temporal unpredictability. However, such a scheme actually leads to a kind of predictability in the sense that, as time continues to pass, the event becomes increasingly likely to occur. Thus, the observer can use time elapsed thus far as a cue to update their expectation that the event is about to occur. Ideally, for maximal unpredictability, time elapsed thus far should not be informative about the impending occurrence of the event.

The evolution of temporal uncertainty about when an event will occur can be more formally characterized by the *hazard rate* (also sometimes called the *hazard function*). The hazard rate measures the ratio of the probability density (for continuous variables) or probability (for discrete variables) that the event will occur at time t to the probability that the event will occur after t. Thus, roughly speaking, the hazard rate measures the odds that an event will occur now or later. For instance, in the example considered above with uniformly distributed time until event occurrence, it can be shown that the evolution of temporal uncertainty follows the hazard rate $h(t) = 1 / (b - t)$, and so the odds of the event occurring now rather than later are 10 times higher when the current time is 90% of the way to the upper bound (i.e. $t = 0.9b$) than they are at the start of the interval (i.e. at $t = 0$).

In general, for any time-varying hazard rate, time elapsed thus far is informative for updating temporal expectation. Thus, to achieve maximal temporal uncertainty, the hazard rate must be constant over time. As discussed below, it is well known that the exponential distribution exhibits a constant hazard rate equal to its rate parameter λ. However, further nuance is required in considering how to apply this result to the construction of random temporal intervals in psychophysics experiments. Whereas the exponential distribution treats time as a continuous variable, computerized psychophysics experiments present stimuli in discrete screen refresh frames (typically around 60 Hz). This effectively discretizes the time at which events can occur, suggesting that here it may be both more appropriate and more convenient to treat time as a discrete variable. Additionally, whereas the exponential distribution has no temporal upper bound, it may often be more practical and more convenient for experimental design to impose an upper bound on the duration of random temporal intervals.

Thus, here we consider the question of how to achieve constant hazard rates when time is modeled as either a continuous or discrete variable, and when the maximum duration of the random temporal interval is unbounded or bounded. For unbounded time intervals, we show how the exponential and geometric distributions yield constant hazard rates for continuous and discrete time, respectively. For bounded time intervals, we discuss how the imposition of the bound poses complications for controlling hazard rate in both continuous and discrete time, but present a solution for exact control of hazard rate in the discrete time case using a modified form of the geometric distribution. The latter solution is likely to be the most relevant for most psychophysics experiments that seek to control hazard rate.

Results

Constant hazard rate for continuous unbounded time intervals

For continuous time t, the hazard rate for event occurrence $h(t)$ is defined as

$$h(t) = \frac{f(t)}{1 - F(t)} \tag{1}$$

where $f(t)$ is a probability density function describing the continuous distribution of time until event occurrence and $F(t)$ is the cumulative distribution function for $f(t)$. For a given time t, the continuous hazard rate thus measures the ratio of the probability density that the event will happen at t to the probability that the event will happen after t.

Consider the case of an observer in a psychophysics task who has been anticipating that an event will occur since time $t = 0$. Currently, at time t_{now}, the event still has not occurred, and the observer is judging the odds of whether the event is about to occur in the next moment or will occur at some unknown time in the future. This observer is then effectively attempting to estimate the hazard rate $h(t_{now})$.[1]

If the hazard rate is constant, then the odds that the event will happen now vs later is constant for all times t, meaning that the time elapsed thus far gives the observer no information about how likely the event is to occur now vs in the future. In other words, the observer's uncertainty about when the event will occur is constant throughout the temporal interval until the event finally occurs.

When t is a continuous variable with no upper bound, the exponential distribution achieves a constant hazard rate λ, since

$$f_{\exp}(t;\ \lambda) = \lambda e^{-\lambda t} \tag{2}$$

$$F_{\exp}(t;\ \lambda) = 1 - e^{-\lambda t} \tag{3}$$

$$h_{\exp}(t;\ \lambda) = \frac{\lambda e^{-\lambda t}}{1 - (1 - e^{-\lambda t})} = \lambda \tag{4}$$

Example exponential distributions and their corresponding hazard rates are shown in **Figure 1**.

[1] See section "Note on hazard rate and dynamic probability updating" below for further discussion.

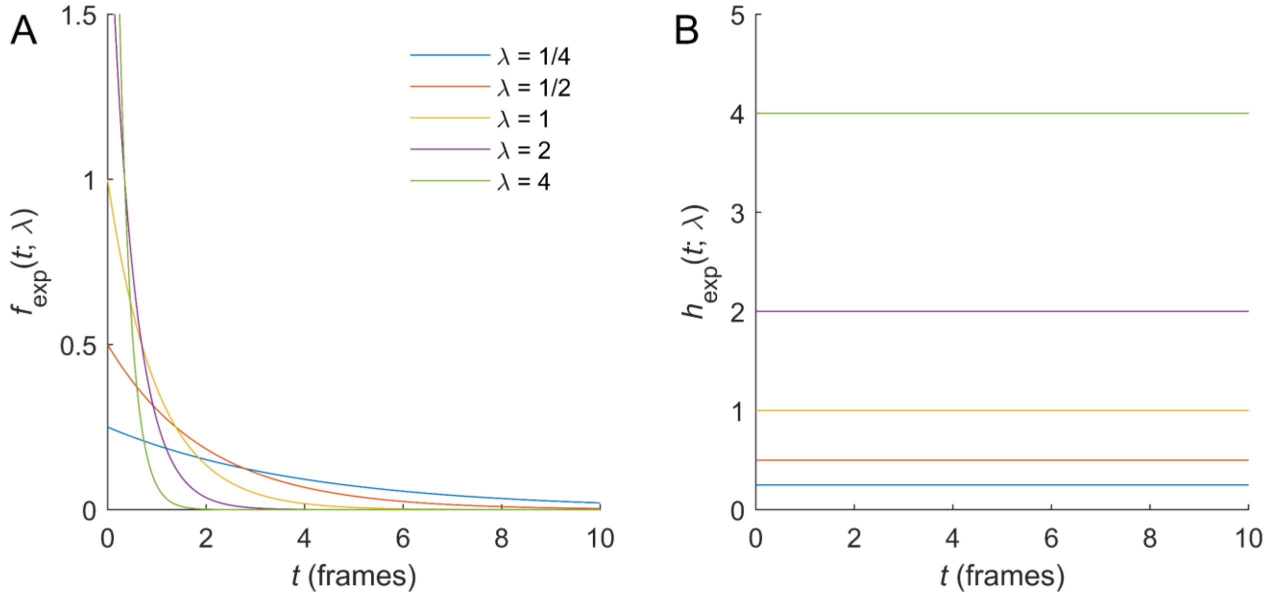

Figure 1. Hazard rates for unbounded exponential distributions. (A) Example exponential distributions with different values of λ. The unit of time is screen refresh frames for consistency with later figures, but note that here time is treated as a continuous variable. The x-axis is cut off at t = 10 frames for display purposes, but the distributions extend to infinity. The y-axis is cut off at 1.5 for display purposes. **(B)** Unbounded exponential distributions have a constant hazard rate equal to λ.

Constant hazard rate for discrete unbounded time intervals

The above result seems to suggest that if one wants to present temporally unpredictable events in a psychophysics task, one can simply arrange for the time of event occurrences relative to some starting time t_0 to be exponentially distributed. However, there is a complicating factor: When stimuli are displayed on a computer screen, the time at which events can occur is discretized by the screen refresh rate. Thus, for experiments conducted on computers, it is more appropriate to model hazard rate with discrete probability functions.

For discrete time n, the hazard rate for event occurrence $h(n)$ is defined as

$$h(n) = \frac{P(N = n)}{1 - P(N \leq n)} \tag{5}$$

where $P(N = n)$ is a probability mass function describing the discrete distribution of time until event occurrence and $P(N \leq n)$ is the cumulative distribution function for $P(N = n)$. For a given time step n, the discrete hazard rate thus measures the ratio of the probability that the event will happen at n to the probability that the event will happen after n.

In the case of an observer anticipating event occurrence on a computer screen, n measures the number of screen refresh frames until the event occurs, starting from the first possible frame at which the event can occur at $n = 1$.

The geometric distribution is the discrete analogue of the exponential distribution and so can be used to achieve a constant hazard rate when n is a discrete variable with no upper bound. It measures the probability that an event occurs at time step n for $n \in \{1, 2, 3, \ldots\}$ when the event can occur at each time step with constant probability p, and thus its PMF is given by

$$P_{\text{geo}}(N = n; p) = (1 - p)^{n-1} p \tag{6}$$

The geometric distribution CDF measures the probability that the event occurs by time step n or earlier, and is thus given by computing the complement of the probability that the event has not yet occurred by time step n for $n \in \{1, 2, 3, \ldots\}$:

$$P_{\text{geo}}(N \leq n; p) = 1 - (1 - p)^{n} \tag{7}$$

The geometric distribution yields a constant hazard rate equal to the odds of event occurrence p / (1 – p), which we will call λ by way of analogy to the constant hazard rate of the exponential distribution:

$$h_{\text{geo}}(n; p) = \frac{(1 - p)^{n-1} p}{1 - (1 - (1 - p)^{n})} = \frac{p}{1 - p} = \lambda \tag{8}$$

The value of p that yields a given hazard rate λ is given by

$$p = \frac{\lambda}{\lambda + 1} \tag{9}$$

Example geometric distributions and their corresponding hazard rates are shown in **Figure 2**.

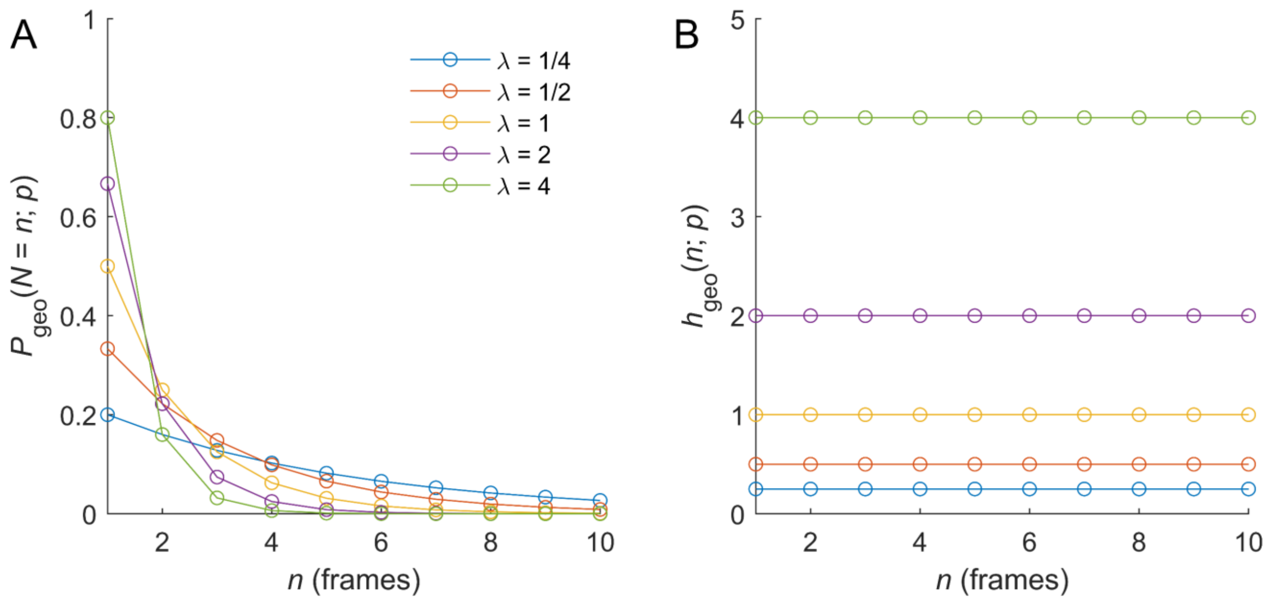

Figure 2. Hazard rates for unbounded geometric distributions. (A) Example geometric distributions with different values of λ. For each distribution, the probability of event occurrence *p* on any given frame is given by $p = \lambda / (\lambda+1)$. Note that here, time is a discrete variable measured in number of screen refresh frames. The x-axis is cut off at $n = 10$ frames for display purposes, but the distributions extend to infinity. **(B)** Unbounded geometric distributions have a constant hazard rate equal to $\lambda = p / (1 - p)$.

Complications for bounded time intervals

The above approaches have allowed for the time until event occurrence to be unlimited. However, in practice, it may be inconvenient to allow for random time intervals with no upper bound in psychophysics experiments. This can occasionally yield very long intervals that may interfere with the participant's engagement in the task and waste precious time resources. More generally, imposing no upper bound on interval length can introduce complications for precisely controlling various other temporal aspects of the experiment, such as the temporal relationships among events within and across trials and total experiment duration. Such considerations can be especially important in neuroimaging contexts.

One way to adjust the above approaches to accommodate an upper bound on interval duration is to use truncated forms of the exponential and geometric distributions. For any probability distribution, truncation at an upper bound *b* is achieved by setting all probability values above *b* to 0 and then normalizing the remaining portions of the distribution so that their cumulative probability is 1.

More formally, the PDF and CDF for the exponential distribution truncated at *b* are given by

$$f_{\text{tr exp}}(t;\ \lambda, b) = \frac{f_{\text{exp}}(t;\ \lambda)}{F_{\text{exp}}(b;\ \lambda)} = \frac{\lambda e^{-\lambda t}}{1 - e^{-\lambda b}}, \qquad t \in [0, b] \tag{10}$$

$$F_{\mathrm{tr\,exp}}(t;\ \lambda, b) = \frac{F_{\exp}(t;\ \lambda)}{F_{\exp}(b;\ \lambda)} = \frac{1 - e^{-\lambda t}}{1 - e^{-\lambda b}}, \qquad t \in [0, b] \tag{11}$$

The expression for the hazard rate of the truncated exponential distribution evaluates to

$$h_{\mathrm{tr\,exp}}(t;\ \lambda, b) = \frac{\lambda}{1 - e^{-\lambda(b-t)}}, \qquad t \in [0, \mathrm{b}) \tag{12}$$

Note that the truncated exponential distribution's hazard rate is undefined at $t = b$ due to the denominator being zero. This is intuitive in the sense that it is impossible for the event to occur after b by definition of the bound, which renders estimation of the odds that the event will occur now or later moot. At $t = b$, the event must occur now and cannot occur later.

Importantly, the truncation procedure has the effect that the hazard rate for the truncated exponential distribution is no longer constant, but rather depends on t and b (Eq. 12). As t approaches b, the hazard rate approaches infinity. Example truncated exponential distributions and their corresponding hazard rates are shown in **Figure 3**.

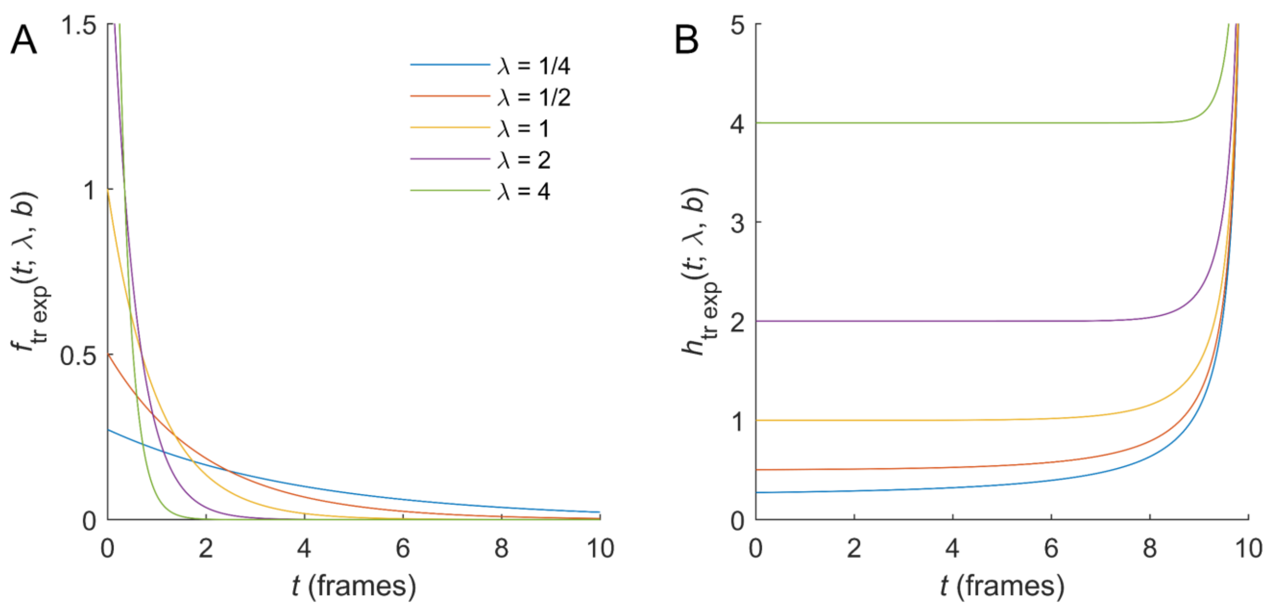

Figure 3. Hazard rates for truncated exponential distributions. (A) Example truncated exponential distributions with different values of λ, truncated at b = 10 frames. The y-axis is cut off at 1.5 for display purposes. **(B)** Truncated exponential distributions have time-varying hazard rates that asymptote to infinity as t approaches b (Eq. 12). By definition of the bound, there is zero probability of the event occurring after b, which entails that the hazard rate is undefined at $t = b$.

A similar phenomenon occurs for the truncated geometric distribution. The PMF and CDF for the truncated geometric distribution are given by

$$P_{\text{tr geo}}(N = n;\ p, b) = \frac{P_{\text{geo}}(N = n;\ p)}{P_{\text{geo}}(N \leq b; p)} = \frac{(1-p)^{n-1}p}{1-(1-p)^b}, \qquad n \in \{1, 2, \dots, b\} \tag{13}$$

$$P_{\text{tr geo}}(N \leq n;\ p, b) = \frac{P_{\text{geo}}(N \leq n;\ p)}{P_{\text{geo}}(N \leq b; p)} = \frac{1-(1-p)^{n}}{1-(1-p)^b}, \qquad n \in \{1, 2, \dots, b\} \tag{14}$$

The expression for the hazard rate of the truncated geometric distribution evaluates to

$$h_{\text{tr geo}}(n;\ p, b) = \frac{p}{(1-p)-(1-p)^{b-n+1}}, \qquad n \in \{1, 2, \dots, b-1\} \tag{15}$$

Note that the truncated geometric distribution's hazard rate is undefined at $n = b$ due to the denominator being zero, as seen above with the truncated exponential distribution.

Importantly, the truncation procedure has the effect that the hazard rate for the truncated geometric distribution is no longer constant, but rather depends on n and b (Eq. 15). Similarly to the behavior of the truncated exponential distribution (**Figure 3**), the hazard rate of the truncated geometric distribution increases with n and becomes undefined at $n = b$. Example truncated geometric distributions and their corresponding hazard rates are shown in **Figure 4**.

Note that the example distributions in **Figure 4** set $b = 10$ frames so that discrete time steps for all times in the bounded interval $n \in [1, b]$ are clearly discernable, and the time scales for all other figures are set similarly for ease of comparison. However, on a standard 60 Hz display, the duration of 10 frames is 167 ms; in practice, the upper bound for time until event occurrence for most experimental use cases would be much larger than this. Nonetheless, for any set of probability distributions resembling those shown in **Figure 4A** over some interval $n \in [1, b]$ (which would require appropriately adjusting the value of the distribution parameter p depending on the value of b), the resulting hazard rates over $n \in [1, b]$ would resemble those shown in **Figure 4B**. Thus, this pattern of results is not specific to the display choice of $b = 10$ frames.

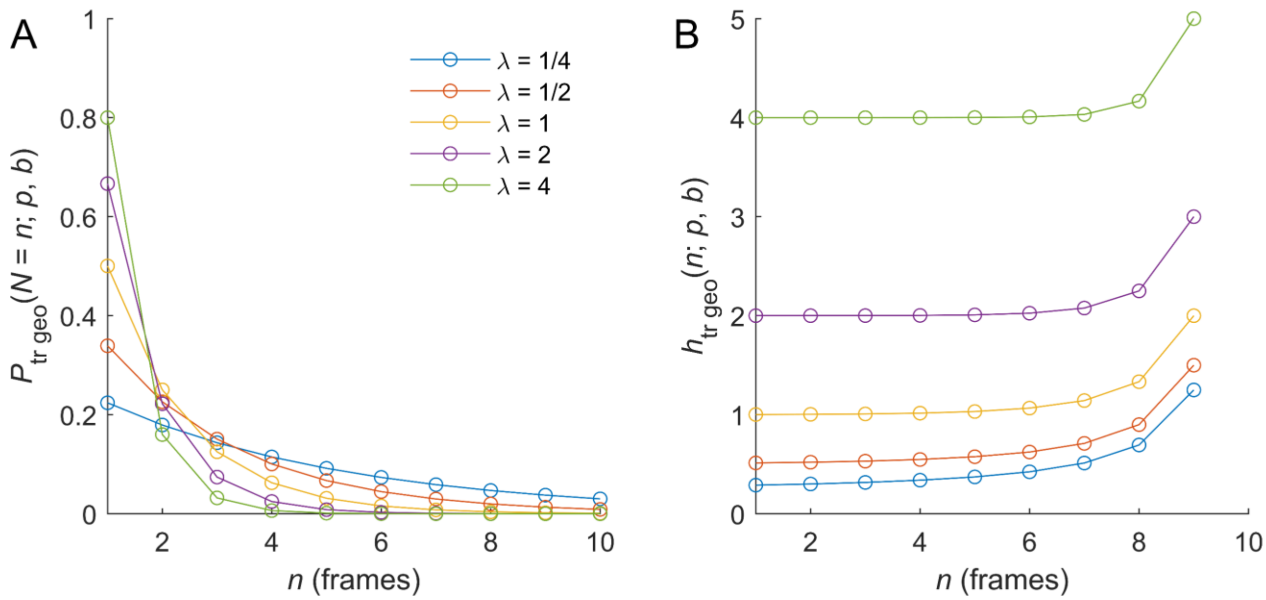

Figure 4. Hazard rates for truncated geometric distributions. (A) Example truncated geometric distributions with different values of λ, truncated at *b* = 10 frames. For each distribution, the probability of event occurrence *p* on any given frame (prior to truncation) is set to λ / (λ+1) (Eq. 9). **(B)** Truncated geometric distributions have time-varying hazard rates that increase with *n* (Eq. 15). By definition of the bound, at *n* = *b* the probability of the event occurring now vs later are 1 and 0, respectively, and so the hazard rate is undefined.

Constant hazard rate for discrete bounded time intervals

The foregoing naturally invites the question of whether it is possible to define alternative probability distributions that exhibit constant hazard rates over bounded time intervals. However, achieving this goal for probability density functions defined over continuous time seems problematic. It is necessarily the case for any time interval bounded at *b* that the hazard rate at *b* is infinite, since the probability of the event occurring after *b* (i.e. the denominator of the hazard rate) is zero by definition. Any continuous probability density function would then necessarily lead to a continuously increasing hazard rate as *t* approaches *b*.

Fortunately, a simple and effective strategy is available in the discrete time case. Above, we considered the behavior of the geometric distribution under the standard definition of distribution truncation in which the PMF is normalized by the CDF over the truncated interval. A natural alternative way to approach truncating the geometric distribution is to leave all probabilities up to the *b* – 1 time step untouched, and to assign all remaining probability to *b*. More formally,

$$P_{\text{cmp geo}}(N = n;\ p, b) = \begin{cases} P_{\text{geo}}(N = n;\ p), & n < b \\ 1 - P_{\text{geo}}(N \le b - 1; p), & n = b \end{cases} \tag{16}$$

We call this the compressed geometric distribution to distinguish it from the truncated geometric distribution defined above[2]. It is “compressed” in the sense that all probability mass in the geometric distribution for $n \geq b$ is compressed into the time step $n = b$. The definition of the compressed geometric distribution in Eq. 16 evaluates to

$$P_{\text{cmp geo}}(N = n;\ p, b) = \begin{cases} (1-p)^{n-1}p, & n < b \\ (1-p)^{b-1}, & n = b \end{cases} \tag{17}$$

The expression at $n = b$ can be interpreted as the probability of the event not occurring for the first $b - 1$ time steps ($(1 - p)^{b-1}$) multiplied by the probability p_b of the event occurring at time step b, where by definition of the bounded time interval, $p_b = 1$.

The CDF and hazard rate for the compressed geometric distribution are given by

$$P_{\text{cmp geo}}(N \leq n;\ p, b) = \begin{cases} (1-p)^{n}, & n < b \\ 1, & n = b \end{cases} \tag{18}$$

$$h_{\text{cmp geo}}(n;\ p, b) = \frac{p}{(1-p)} = \lambda, \qquad n \in \{1, 2, \ldots, b-1\} \tag{19}$$

The hazard rate for the compressed geometric distribution (Eq. 17) thus exhibits the desired behavior: It yields the same constant hazard rate as the standard geometric distribution (Eq. 8) for all time steps $n < b$ where the hazard rate is defined. This hazard rate is necessarily undefined at $n = b$ due to how the general definition of the hazard rate breaks down at the upper bound of a bounded temporal interval, as considered in the above discussion on truncated distributions. Example compressed geometric distributions and their corresponding hazard rates are shown in **Figure 5**.

[2] This approach has formal similarities to the construction of probability distributions from censored data, but conceptually is not related to data censoring. Thus, here we use the term “compressed distribution” rather than “censored distribution” for conceptual clarity.

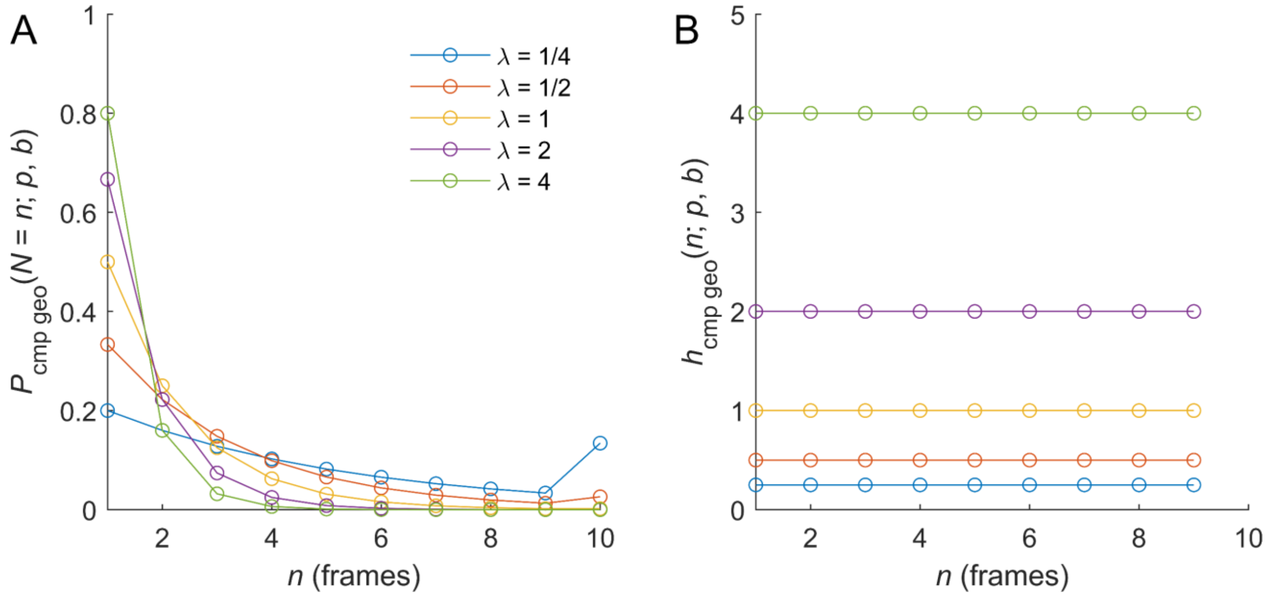

Figure 5. Hazard rates for compressed geometric distributions. (A) Example compressed geometric distributions with different values of λ, with all probability mass for $n > b$ frames in the corresponding standard geometric distributions compressed into the $n = b$ frame, for $b = 10$. For each distribution, the probability of event occurrence p for all frames $n < b$ is set to $\lambda / (\lambda+1)$ (Eq. 9). By definition of the bound, probability of event occurrence at frame b is 1 (Eq. 17). **(B)** Compressed geometric distributions have constant hazard rates for all time steps n where the hazard rate is defined. As with truncated geometric distributions (Figure 4), hazard rate is necessarily undefined at the final possible frame for event occurrence b.

Note on hazard rate and dynamic probability updating

In the above discussion, we have framed the hazard rate as roughly indicating the odds that an event will happen now (at $t = t_{now}$) vs later (at $t > t_{now}$). However, our formal treatment of hazard rate has assumed a static probabilistic knowledge of when the event will occur for all time values starting from $t = 0$ (or $n = 1$ in the discrete case). But if the event is very easily perceived when it does occur, then an attentive observer still waiting for event occurrence at time $t_{now} > 0$ would know that the event has not yet occurred, i.e. they would know that the probability that the event occurred at some previous time $0 \leq t < t_{now}$ is zero.[3] It follows that the observer's probability estimates of event occurrence for all previous times $t < t_{now}$ should be updated from their initial values to zero to reflect their new knowledge about event occurrence at t_{now}. In light of this point, is it really valid to characterize the epistemic situation of the observer at t_{now} with the hazard rate evaluated at $h(t_{now})$ when $t_{now} > 0$, given that our treatment of $h(t_{now})$ thus far does not feature any updating of event occurrence probabilities as time passes?

[3] If the event is presented at threshold levels of detection, such that it is not always detected when it does occur, then dynamic updating of the probability of event occurrence for previous times $t < t_{now}$ is considerably more complicated. For simplicity, here we consider only the suprathreshold case and leave consideration of the threshold case to future work.

Dynamically updating probability estimates to zero for all previous time steps $t < t_{\mathrm{now}}$ and renormalizing all remaining probabilities so they sum to 1 is a truncation operation on the lower portion of the probability distribution, similar to the truncation operation on the upper portions of distributions we used above to impose an upper bound b on time intervals. In the most general treatment, for a continuous PDF $f(t)$ and its corresponding CDF $F(t)$ and hazard rate $h(t)$, updated formulae after truncating the lower and upper portions of $f(t)$ at t_{now} and b are given by

$$f_{\mathrm{tr}}^{\mathrm{now}}(t;\ t_{\mathrm{now}}, b) = \frac{f(t)}{F(b) - F(t_{\mathrm{now}})}, \qquad t \in [t_{\mathrm{now}}, b] \tag{20}$$

$$F_{\mathrm{tr}}^{\mathrm{now}}(t;\ t_{\mathrm{now}}, b) = \frac{F(t) - F(t_{\mathrm{now}})}{F(b) - F(t_{\mathrm{now}})}, \qquad t \in [t_{\mathrm{now}}, b] \tag{21}$$

$$h_{\mathrm{tr}}^{\mathrm{now}}(t;\ t_{\mathrm{now}}, b) = \frac{f_{\mathrm{tr}}^{\mathrm{now}}(t;\ t_{\mathrm{now}}, b)}{1 - F_{\mathrm{tr}}^{\mathrm{now}}(t;\ t_{\mathrm{now}}, b)} = \frac{f(t)}{F(b) - F(t)}, \qquad t \in [t_{\mathrm{now}}, b) \tag{22}$$

Truncating the lower portion of the PDF and CDF at t_{now} is achieved by the $F(t_{\mathrm{now}})$ terms in Eq. 20 and 21. However, these terms cancel out in the formula for hazard rate, as reflected by the absence of any $F(t_{\mathrm{now}})$ terms in the most simplified form of Eq. 22. Additionally, when no upper bound is imposed (i.e. $b = \infty$), Eq. 22 reduces to the equation for hazard rate over unbounded continuous time given in Eq. 1. These conclusions hold not just for the exponential distribution, but for any PDF $f(t)$. Similar conclusions can readily be drawn for the discrete time case.

Thus, truncating the probability distribution for event occurrence at t_{now} does not influence the hazard rate $h(t)$ for $t \geq t_{\mathrm{now}}$. It follows that our characterization of the observer using a static form of $h(t)$ to estimate the odds of the event occurring now or later at t_{now} is equivalent to the scenario where the event is presented at easily perceived suprathreshold levels, and so the observer dynamically updates probability estimates of event occurrence as time passes by setting probability for all previous times $t < t_{\mathrm{now}}$ to zero and renormalizing all remaining probabilities before estimating an updated $h(t)$.

Discussion

In the foregoing we considered how best to achieve the goal of making the time at which some event occurs in a given trial of a psychophysics experiment maximally uncertain, with the intention of minimizing possible effects of expectation and temporal attention on behavioral and neural data. Assuming the observer is shown many trials in which the event occurs, and that across trials the time elapsed until event occurrence follows some probability distribution, what choice of distribution best achieves the goal of making the event's occurrence most unpredictable?

We framed this question in terms of finding the probability distribution that yields a constant hazard rate $h(t)$ for event occurrence over time. Hazard rate measures the ratio of probability (or probability density for PDFs) that the event will occur at time t to the probability that the event will occur after t. Thus, we can think of hazard rate as measuring the odds that the event will happen now or later. When hazard rate is constant, time elapsed thus far is not informative about the impending occurrence of the event, and so the observer's uncertainty is constant even as time continues to pass.

We first reviewed how the exponential distribution achieves a constant hazard rate when time is continuous and unbounded, but we noted that applying this result to an experimental design requires some further nuance. Computerized tasks displaying stimuli on screens discretize the time at which events can occur, suggesting that a discrete probability distribution is both more appropriate and convenient to use in this case. Additionally, we noted that to precisely control various other temporal aspects of an experiment, it may be desirable to impose an upper bound on the duration of random temporal intervals.

We showed how the geometric distribution, a natural discrete analogue of the continuous exponential distribution, yields constant hazard rates when time is unbounded. However, we also showed that both the exponential and geometric distributions no longer yield constant hazard rates when they are truncated to prevent durations greater than some upper bound b using the standard definition of distribution truncation. In both cases, the hazard rate increases as time approaches the upper bound. At $t = b$, the definition of the hazard rate breaks down, since by definition of the bound the event has zero probability of occurring after b. Mathematically, this leads to a division by zero; conceptually, it corresponds to the fact that there is no point in estimating the odds that an event will happen now or later when it is known for a fact that it can't occur later.

However, in the discrete time case, we show that an alternative way of truncating the geometric distribution yields the desired behavior of exhibiting constant hazard rate for all times $n < b$ where the hazard rate is defined. The standard approach to distribution truncation over an interval [a, b] involves setting all probabilities outside that interval to zero, and then normalizing the remaining probabilities in the interval so that they sum to one. However, an alternative approach for setting an upper bound on a discrete probability distribution is to assign the cumulative probability for time steps $n \geq b$ to time step b, and set all probabilities above b to zero. This approach has formal similarities with the construction of probability distributions from censored data, but conceptually is not related to data censoring. Thus, for conceptual clarity, here we use the term "compressed distribution" to refer to the resulting distribution and to differentiate it from the standard definition of the truncated distribution. We showed that the compressed geometric distribution exhibits constant hazard rate for all $n < b$ where hazard rate is defined. Hazard rate remains undefined at $n = b$, as must be the case for any hazard rate over a bounded time interval.

Thus, for experimenters seeking to minimize the predictability of events that occur over many trials, we recommend achieving this goal by arranging for the time until event occurrence to follow the geometric distribution, where discrete time steps correspond to screen refresh frames. If the number of frames until event occurrence n_{event} does not need to be bounded, then n_{event} on any given trial can be randomly drawn from the geometric distribution $P_{\text{geo}}(N = n; p)$ (Eq. 6). Alternatively, if the number of frames until event occurrence is bounded at b, then n_{event} on any given trial can be randomly drawn from the compressed geometric distribution $P_{\text{cmp geo}}(N = n;\ p, b)$ (Eq. 16). The average duration of the random interval can be controlled with the parameter *p*, which measures the probability of the event occurring on any given frame.

Acknowledgements

We thank Rachel Denison, Karen Tian, Michael Epstein, and Angela Shen for helpful discussions in the development of this methodology. We also acknowledge support from the Templeton World Charity Foundation ("An Adversarial Collaboration to Test Predictions of First-Order and Higher-Order Theories of Consciousness", TWCF Number: 0567, to M.A.K.P. and R.D.) and the Canadian Institute for Advanced Research (Fellowship in the Brain, Mind, & Consciousness Program, to M.A.K.P.). The funders had no role in the design or execution of this project.