

Memory Nearly on a Spring: A Mean First Passage Time Approach to Memory Lifetimes

Terry Elliott¹

Department of Electronics and Computer Science,

University of Southampton,

Highfield,

Southampton, SO17 1BJ,

United Kingdom.

Running Title: Memory on a Spring

March 13, 2014

¹Tel.: +44 (0)23 8059 6000, Fax.: +44 (0)23 8059 2783, E.-mail:
te@ecs.soton.ac.uk.

Abstract

We study memory lifetimes in a perceptron-based framework with binary synapses, using the mean first passage time for the perceptron's total input to fall below firing threshold to define memory lifetimes. Working with the simplest, memory-related model of synaptic plasticity, we may obtain exact results for memory lifetimes or, working in the continuum limit, good analytical approximations that afford either much qualitative insight or extremely good quantitative agreement. In one particular limit, we find that memory dynamics reduce to the well-understood Ornstein-Uhlenbeck process. We show that asymptotically, the lifetimes of memories grow logarithmically in the number of synapses when the perceptron's firing threshold is zero, reproducing standard results from signal-to-noise ratio analyses. However, this is only an asymptotically valid result, and we show that extending its application outside the range of its validity leads to a massive over-estimate of the minimum number of synapses required for successful memory encoding. In the case that the perceptron's firing threshold is positive, we find the remarkable result that memory lifetimes are strictly bounded from above. Asymptotically, the dependence of memory lifetimes on the number of synapses drops out entirely, and this asymptotic result provides a strict upper bound on memory lifetimes away from this asymptotic regime. The classic, logarithmic growth of memory lifetimes in the simplest, palimpsest memories is therefore untypical and non-generic: memory lifetimes are typically strictly bounded from above.

1 Introduction

Accumulating evidence suggests that synaptic strengths may take only a finite, discrete set of values (Petersen et al., 1998; Montgomery and Madison, 2002, 2004; O'Connor et al., 2005a, 2005b, 2007) or change only in discrete steps (Yasuda et al., 2003; Bagal et al., 2005; Sobczyk and Svoboda, 2007); the former implies the latter although the converse need not follow. Imposing biophysically-motivated bounds on synaptic strengths or employing discrete-strength, typically binary-strength synapses in mathematical models of memory formation overcomes the catastrophic forgetting associated with the standard Hopfield network (Hopfield, 1982; Hertz et al., 1991), instead turning memory systems into palimpsests, which learn new memories by forgetting old ones (Nadal et al., 1986; Parisi, 1986; Tsodyks, 1990; Amit and Fusi, 1994). Such models, however, run straight onto the horns of the plasticity-versus-stability dilemma (Abraham and Robins, 2005), with the fidelity of the recall of synaptically-stored memories in the simplest models typically falling exponentially fast as changes in synapses' strengths induced by ongoing memory storage degrade already-stored memories.

The determination of memory lifetimes in feed-forward or recurrent networks with binary-strength synapses has become an area of considerable interest because excessively short memory lifetimes may invalidate models as viable candidates for at least long-term memory storage, either in real brains or in electronic hardware. Critical to this determination is the very definition of the

lifetime of a memory. The simplest and perhaps by-now classic definition of memory lifetime is based on a signal-to-noise ratio (SNR) analysis of a memory (Tsodyks, 1990; Amit and Fusi, 1994; Brader et al., 2007) although alternative approaches exist, including for example signal detection theory (Leibold and Kempter, 2006, 2008) and retrieval probabilities (Huang and Amit, 2010, 2011). In an SNR analysis in a feed-forward, perceptron system, the ratio of the mean (unthresholded) output of a neuron to its standard deviation in response to the memory of interest is used to define that memory’s lifetime: the point at which the SNR drops to unity is taken to define the point in time at which the memory has been forgotten. The dependence of this lifetime on the perceptron’s number of synapses, N , essentially determines the efficiency of the memory system. For a palimpsest in the simplest models, this dependence is only logarithmic (Tsodyks, 1990).

SNRs may, however, be problematic for a variety of reasons. First, there is a degree of arbitrariness in setting the SNR threshold of unity. Second, by considering only the variance as a source of fluctuations that can help to drive the SNR below unity, potentially important higher-order fluctuations are ignored. Third, SNR memory lifetimes differ depending on whether we consider a discrete-time formulation or a continuous-time formulation for memory storage (Elliott and Lagogiannis, 2012). Finally, the SNR is somewhat of an abstraction divorced from a neuron’s firing characteristics: as a quantity involving statistics of different orders, it is hard to relate an SNR directly to any particular property of a system.

Here, as an alternative to SNRs, we consider defining memory lifetimes by the mean first passage time (MFPT) for a perceptron's (unthresholded) output to drop below firing threshold. This definition is free from arbitrariness; it includes all possible sources of fluctuations; MFPTs, as we shall show, are identical in discrete-time and continuous-time formulations. Perhaps more important, the first passage time for a perceptron's total input to fall below firing threshold is a quantity of direct and immediate relevance to the neuron: it can, for example, be regarded as a direct read-out of membrane integrative dynamics, determining the first point in time at which re-presentation of a stored memory will fail to induce the perceptron to fire, thus defining the point at which the neuron has forgotten the memory. The *mean* first passage time is then a natural quantity to consider, averaging over many realisations but without mixing up statistics of different orders. By using an MFPT definition of memory lifetimes in the simplest, memory-related model of synaptic plasticity, we show that for zero firing threshold, we indeed obtain asymptotic logarithmic dependence of memory lifetimes on the number of synapses. However, this is only an asymptotic result and we show that it typically massively over-estimates the minimum number of synapses required for successful memory storage. For a non-zero firing threshold, we find the remarkable result that even this rather feeble logarithmic asymptotic behaviour is broken. Specifically, we find that for any positive firing threshold, for N large enough, memory lifetimes asymptote to a *constant*, independent of N . For positive firing thresholds, therefore, memory lifetimes are strictly bounded from above

for all N .

The remainder of this paper is organised as follows. In section 2 we set up the problem and establish our basic notation. In the following section, we derive generic results for MFPT memory lifetimes, for both discrete-time and continuous-time formulations. In section 4 we introduce and define the simplest possible, memory-related model of synaptic plasticity (Tsodyks, 1990). Next, restricting to this specific model, we compute exact transition probabilities and MFPTs. In section 6 we take the continuum limit of this model, in which the perceptron's output is regarded as a continuous- rather than discrete-valued variable. Doing so enables us to develop an approximation that provides much qualitative understanding for comparatively little analytical effort. In section 7, we formulate continuum MFPTs in terms of an integral equation involving jump processes and consider two methods for solving this equation by making approximations to the kernel. We then compare analytical, numerical and simulation results in section 8. Finally we discuss our approach and results and consider future possible extensions to more complicated models of synaptic plasticity.

2 Preliminaries

Consider a single neuron or perceptron with N synapses of strengths $S_i(t) \in \{-1, +1\}$, $i = 1, \dots, N$, at time $t \geq 0$. With inputs x_i , $i = 1, \dots, N$, through

these N synapses, the neuron's output is taken to be of the standard form

$$h_{\mathbf{x}}(t) = \frac{1}{N} \sum_{i=1}^N x_i S_i(t), \quad (2.1)$$

where the factor $1/N$ is a normalisation term that may be dropped under a rescaling of the neuron's firing threshold. For simplicity, we do not consider a bias because it may also be absorbed into a re-definition of the neuron's firing threshold. The inputs x_i are assumed without loss of generality to have two levels, $x_i \in \{-1, +1\}$, with $x_i = -1$ corresponding to low input and $x_i = +1$ corresponding to high input (see Hertz et al., 1991). The two strength states ± 1 for a synapse are interpreted as weak (-1) and strong ($+1$) rather than inhibitory and excitatory, without loss of generality. We may, for example, write $S'_i(t) = \frac{1}{2}[1 + S_i(t)] \in \{0, 1\}$ and work with $S'_i(t)$ instead of $S_i(t)$. A bias of the form $-\frac{1}{N} \sum_{i=1}^N x_i$ would then be induced. When the inputs x_i have zero mean and unit variance, this induced bias has zero mean and variance $1/N$ and so is essentially irrelevant; it may, in any event, be absorbed into a re-definition of the neuron's firing threshold. The neuron's actual output is set to $+1$ or -1 according to whether $h_{\mathbf{x}}(t) > \theta$ or $h_{\mathbf{x}}(t) \leq \theta$, respectively, although we do not explicitly threshold the output in the following and we will simply refer to $h_{\mathbf{x}}(t)$ in Eq. (2.1) as the neuron's output, or sometimes, when we draw the distinction, its total input.

The neuron is required to store a set of "synaptic memories" $\boldsymbol{\xi}^\alpha$, where $\alpha = 0, 1, 2, \dots$ indexes these memories. Any particular memory $\boldsymbol{\xi}^\alpha$ constitutes an input \mathbf{x} to the neuron at some defined time and induces changes in the synaptic

strengths that lead to the storage of ξ^α . The memories ξ^α have components $\xi_i^\alpha = \pm 1$ and these are assumed for simplicity to be independent both over $i = 1, \dots, N$ and across the different memories $\alpha = 0, 1, 2, \dots$. We need not, for a single perceptron, specify a target neuronal output $h^\alpha \in \{-1, +1\}$ for each memory input vector $+\xi^\alpha$ since for $h^\alpha = -1$ we may without loss of generality reverse the sign of the input vector, to $-\xi^\alpha$, instead. The target outputs are therefore always $h^\alpha = +1$. With this convention, the component ξ_i^α is the synaptic plasticity induction signal to synapse i upon presentation of memory α , with $\xi_i^\alpha = +1$ corresponding to a strengthening or potentiating induction signal and $\xi_i^\alpha = -1$ corresponding to a weakening or depressing induction signal.

For binary-valued synapses, the neuron is a palimpsest, storing newer memories by forgetting older ones (cf. Nadal et al., 1986; Parisi, 1986), overcoming the catastrophic forgetting associated with unbounded, continuous-valued synapses in the Hopfield network (Hopfield, 1982; Hertz et al., 1991). The lifetime of a memory may be gauged by the period for which the memory can be successfully retrieved after its initial storage. Often this is defined via an SNR, a ratio of the neuron's mean output to the standard deviation in the neuron's output. Neither quantity is available to the neuron itself and so an SNR is somewhat of an abstraction. However, the neuron's (unthresholded) output $h_{\xi^\alpha}(t)$ for any particular memory ξ^α is directly available, and we may instead define a memory's lifetime as the average duration for which $h_{\xi^\alpha}(t) > \theta$ after storage, i.e. as the MFPT for $h_{\xi^\alpha}(t)$ to fall below firing threshold. Such

an approach has the advantage of considering all possible fluctuations leading to the neuron’s output falling below threshold and not just the fluctuations due to the variance.

In a discrete-time framework, memory α is stored at discrete time step α , while in a continuous-time framework, memories are stored at times determined by a Poisson process of rate r , although the first memory, ξ^0 , is always stored at time $t = 0$ s. In either case, we track the particular memory ξ^0 stored at $t = 0$ s and determine its lifetime; we refer to this memory as the tracked memory and the neuron’s subsequent output $h_{\xi^0}(t)$ as the tracked memory signal, or just the memory signal. We write $h(t) = h_{\xi^0}(t)$ for simplicity. Since we are not interested in the lifetime of any particular realisation of a pattern ξ^0 , we will typically average over all possible initial patterns ξ^0 .

Model synapses may or may not possess internal, plasticity-related states, where any internal states may be modified by plasticity induction signals and may affect the subsequent expression of plasticity by synapses. For generality, in this and in the following section we consider that a synapse may possess such internal states, although in later sections we will restrict for analytical simplicity to synapses without internal structure. Let $\mathbf{S}_i(t)$ be a vector that denotes the joint probability distribution of both the strength of synapse i and its current internal state, at time t . We will assume that any internal states are symmetric between weak and strong states, and that the first set of, say, s entries in $\mathbf{S}_i(t)$ correspond to the distribution of internal states when the synapse is weak and the second set of s entries to the distribution when the

synapse is strong. Writing

$$\mathbf{n}_- = (\underbrace{1, \dots, 1}_s | \underbrace{0, \dots, 0}_s)^T, \quad (2.2)$$

$$\mathbf{n}_+ = (\underbrace{0, \dots, 0}_s | \underbrace{1, \dots, 1}_s)^T, \quad (2.3)$$

where a superscript T denotes the transpose, the probability that synapse i is strong or weak at time t is then simply

$$\text{Prob}[S_i(t) = \pm 1] = \mathbf{n}_\pm^T \mathbf{S}_i(t), \quad (2.4)$$

i.e. we just sum over the relevant part of the joint probability distribution to integrate out the internal states.

Let \mathbb{M}_+ be the $2s \times 2s$ matrix that represents the transitions in state $\mathbf{S}_i(t)$ when the synapse is subjected to a potentiating induction signal, and \mathbb{M}_- that when subjected to a depressing induction signal. For the storage of $\boldsymbol{\xi}^0$, the synapses are subjected to definite inductions signals, either +1 or -1, acting on the asymptotic or equilibrium distribution for $\mathbf{S}_i(t)$, which we shall discuss later for a specific model of synaptic plasticity. For subsequent memory storage, however, ξ_i^α , $\alpha > 0$, may be +1 with probability 1/2 or -1 with probability 1/2. The state transition matrix for the storage of each subsequent memory at any synapse is then just

$$\mathbb{M} = \frac{1}{2} (\mathbb{M}_+ + \mathbb{M}_-). \quad (2.5)$$

Changes in strength and internal state upon the presentation of memory α are assumed to be independent across all synapses, i.e. synapse i responds only to

its own local induction signal ξ_i^α . Thus, the operator that represents the simultaneous updating of all N synapses upon subsequent memory presentation is the tensor product operator

$$\mathbb{M}_N = \underbrace{\mathbb{M} \otimes \cdots \otimes \mathbb{M}}_N. \quad (2.6)$$

This operator acts on the product space of states of each synapse, which we may write as $\mathbf{S}(t) = \mathbf{S}_1(t) \otimes \cdots \otimes \mathbf{S}_N(t)$, for a definite ordering of synapses.

3 Mean First Passage Time for Perceptron

Output

With these preliminaries completed, we may now derive an expression for the MFPT for $h_{\xi^0}(t)$ to fall below θ , which we denote by $\tau_{\text{mfpt}}(\xi^0)$, in both discrete time and continuous time. The MFPTs for discrete-time and continuous-time processes are, up to a simple multiplicative rate factor, identical, unlike memory lifetimes obtained from SNR calculations. To derive an expression for the MFPT, we must project out of \mathbb{M}_N any transitions that lead to $h_{\xi^0} \leq \theta$, or retain only those transitions maintaining $h_{\xi^0} > \theta$. Let

$$\mathbb{P}_- = \left(\begin{array}{c|c} \mathbb{I}_s & \mathbb{O}_s \\ \hline \mathbb{O}_s & \mathbb{O}_s \end{array} \right) \quad \text{and} \quad \mathbb{P}_+ = \left(\begin{array}{c|c} \mathbb{O}_s & \mathbb{O}_s \\ \hline \mathbb{O}_s & \mathbb{I}_s \end{array} \right), \quad (3.1)$$

with these matrices acting on $\mathbf{S}_i(t)$, projecting out only weak states (\mathbb{P}_-) and strong states (\mathbb{P}_+), where \mathbb{I}_s is the $s \times s$ identity matrix and \mathbb{O}_s is the $s \times s$ zero matrix. These are projection matrices satisfying $(\mathbb{P}_+)^2 = \mathbb{P}_+$, $(\mathbb{P}_-)^2 = \mathbb{P}_-$

and $\mathbb{P}_+\mathbb{P}_- = \mathbb{O}_{2s}$, with \mathbb{O}_{2s} being the $2s \times 2s$ zero matrix. Define

$$\mathbb{P}_{h>\theta} = \sum_{\text{permutations: } h>\theta} \underbrace{\mathbb{P}_{\pm} \otimes \cdots \otimes \mathbb{P}_{\pm}}_N, \quad (3.2)$$

where we sum over that subset of the 2^N possible tensor products that enforce $h_{\xi^0} > \theta$. For example, for $N = 4$ and $\theta = 0$, with say $\xi^0 = (+1, +1, -1, -1)^T$, this operator would be

$$\begin{aligned} \mathbb{P}_{h>0} = & \mathbb{P}_+ \otimes \mathbb{P}_+ \otimes \mathbb{P}_- \otimes \mathbb{P}_- & \leftarrow h \equiv +1 \\ & + \mathbb{P}_- \otimes \mathbb{P}_+ \otimes \mathbb{P}_- \otimes \mathbb{P}_- & \leftarrow h \equiv +1/2 \\ & + \mathbb{P}_+ \otimes \mathbb{P}_- \otimes \mathbb{P}_- \otimes \mathbb{P}_- & \leftarrow h \equiv +1/2 \\ & + \mathbb{P}_+ \otimes \mathbb{P}_+ \otimes \mathbb{P}_+ \otimes \mathbb{P}_- & \leftarrow h \equiv +1/2 \\ & + \mathbb{P}_+ \otimes \mathbb{P}_+ \otimes \mathbb{P}_- \otimes \mathbb{P}_+ & \leftarrow h \equiv +1/2 \end{aligned}$$

with the other $16 - 5 = 11$ permutations projecting onto the forbidden $h_{\xi^0} \leq 0$ states. Because each permutation is distinct in $\mathbb{P}_{h>\theta}$, it is clear that $\mathbb{P}_{h>\theta}$ is also a projection operator, i.e. $(\mathbb{P}_{h>\theta})^2 = \mathbb{P}_{h>\theta}$, because any cross terms in the product vanish due to $\mathbb{P}_+\mathbb{P}_- = \mathbb{O}_{2s}$. The transition operator that implements synaptic updates under memory storage while enforcing only $h_{\xi^0} > \theta$ to $h_{\xi^0} > \theta$ transitions without h_{ξ^0} dropping below threshold is then just $\mathbb{P}_{h>\theta}\mathbb{M}_N$. Let the initial state of all N synapses after the storage of ξ^0 be $\mathcal{P}(\xi^0)$. This may contain states corresponding to $h_{\xi^0} \leq \theta$, and they must be removed too, so the initial, allowed state is just $\mathbb{P}_{h>\theta}\mathcal{P}(\xi^0)$. The state after m subsequent memory storage events is then

$$(\mathbb{P}_{h>\theta}\mathbb{M}_N)^m [\mathbb{P}_{h>\theta}\mathcal{P}(\xi^0)] \equiv (\mathbb{P}_{h>\theta}\mathbb{M}_N\mathbb{P}_{h>\theta})^m [\mathbb{P}_{h>\theta}\mathcal{P}(\xi^0)]. \quad (3.3)$$

where we have used the projection property $(\mathbb{P}_{h>\theta})^2 = \mathbb{P}_{h>\theta}$ to write the $h_{\xi^0} > \theta$ to $h_{\xi^0} > \theta$ transition operator as $\mathbb{P}_{h>\theta}\mathbb{M}_N\mathbb{P}_{h>\theta}$ rather than $\mathbb{P}_{h>\theta}\mathbb{M}_N$ because, as an operator, $\mathbb{P}_{h>\theta}\mathbb{M}_N\mathbb{P}_{h>\theta}$ contains fewer non-zero entries and is thus somewhat easier to handle. With $\mathbf{n} = \mathbf{n}_+ + \mathbf{n}_-$ and writing

$$\mathcal{N} = \underbrace{\mathbf{n} \otimes \cdots \otimes \mathbf{n}}_N, \quad (3.4)$$

the probability that $h_{\xi^0} > \theta$ after m subsequent memory encoding events without ever having dropped below threshold is

$$p_m = \mathcal{N}^T (\mathbb{P}_{h>\theta}\mathbb{M}_N\mathbb{P}_{h>\theta})^m [\mathbb{P}_{h>\theta}\mathcal{P}(\xi^0)]. \quad (3.5)$$

Because $\mathbb{P}_{h>\theta} \neq \mathbb{I}_{2s} \otimes \cdots \otimes \mathbb{I}_{2s}$, $p_m \rightarrow 0$ as $m \rightarrow \infty$, i.e. it is inevitable that h_{ξ^0} will eventually drop below threshold. As a sequence, p_0, p_1, p_2, \dots is monotonic decreasing, and the “leak” $p_m - p_{m+1}$, $m \geq 0$, is the probability that h_{ξ^0} falls below threshold when memory $\alpha = m + 1$ is stored. The MFPT for the discrete-time process is then just

$$\begin{aligned} \tau_{\text{mfpt}}(\xi^0) &= \sum_{m=0}^{\infty} (m+1) (p_m - p_{m+1}) \\ &= \sum_{m=0}^{\infty} p_m \\ &= \mathcal{N}^T \left[\sum_{m=0}^{\infty} (\mathbb{P}_{h>\theta}\mathbb{M}_N\mathbb{P}_{h>\theta})^m \right] [\mathbb{P}_{h>\theta}\mathcal{P}(\xi^0)] \\ &= \mathcal{N}^T \left[(\mathbb{I}_{2s} \otimes \cdots \otimes \mathbb{I}_{2s}) - (\mathbb{P}_{h>\theta}\mathbb{M}_N\mathbb{P}_{h>\theta}) \right]^{-1} [\mathbb{P}_{h>\theta}\mathcal{P}(\xi^0)]. \end{aligned} \quad (3.6)$$

$$(3.7)$$

In continuous time, the memories ξ^α , $\alpha > 0$, are presented at times determined by a Poisson process of rate r . The probability of precisely m memories being

presented in time t is just $e^{-rt}(rt)^m/m!$ and the probability that $h_{\xi^0} > \theta$ after a time t without ever having dropped below threshold is given by the Poisson sum

$$p(t) = \sum_{m=0}^{\infty} \frac{(rt)^m}{m!} e^{-rt} p_m, \quad (3.8)$$

with $-dp(t)/dt \geq 0$ corresponding to the “leak” at time t . For the continuous-time process, the MFPT is then

$$\begin{aligned} \tau_{\text{mfpt}}(\xi^0) &= - \int_0^{\infty} dt t \frac{dp(t)}{dt} \\ &= \int_0^{\infty} dt p(t) \\ &= \sum_{m=0}^{\infty} p_m \int_0^{\infty} dt \frac{(rt)^m}{m!} e^{-rt} \\ &= \frac{1}{r} \sum_{m=0}^{\infty} p_m. \end{aligned} \quad (3.9)$$

Thus, up to the overall factor of the rate r , the MFPTs are identical in both discrete time and continuous time, regardless of the model of synaptic plasticity encoded in the underlying transition matrix \mathbb{M} . This result does not in fact depend on a specifically Poisson process as the continuous-time process. In Appendix A we show that the result carries over to any continuous-time process in which the waiting times between memory storage events are governed by some common probability density function that need not be the exponential distribution.

In principle, in order to determine the MFPT for the neuron’s output in response to memory ξ^0 to fall below firing threshold, we simply have to compute the inverse of the tensor product operator in Eq. (3.7). In practice, however, such a computation is extremely hard if not intractable for all but the most

trivial of problems. For $\theta = 0$, for example, the projection operator $\mathbb{P}_{h>0}$ contains 2^{N-1} terms for N odd and $2^{N-1} - \frac{1}{2} {}^N C_{N/2}$ terms for N even, where ${}^N C_{N/2}$ is a binomial coefficient, and the tensor product operator to be inverted does not factorise over the underlying synaptic space (as we might expect, because $h_{\xi^0}(t)$ is a function over all synapses). Perhaps the simplest strategy is to “flatten” the product space, although the resulting flattened matrices become $(2s)^N \times (2s)^N$, and except for very small N , they will be unmanageably large. Another strategy is to work directly in the transitions in the discrete values of the neuron’s output, which ranges from -1 to $+1$ in steps of $2/N$. Matrices would then be only $(N+1) \times (N+1)$. However, with such a collapsing of the underlying synaptic dynamics down to just h_{ξ^0} transitions, these transitions become non-Markovian except for synapses without internal states. To compute memory lifetimes from a MFPT approach, we must therefore typically develop approximation methods, use numerical methods, or resort to simulations to extract results. Here, we will restrict to a particular model of synaptic plasticity for which essentially exact results may be obtained, or for which good, controlled approximations are available.

We will compare memory lifetimes determined from MFPTs to those from SNRs. With SNRs, a memory is typically deemed to have been forgotten when the SNR reaches a value of unity. At this point, the memory signal becomes indistinguishable from zero at the one standard deviation level. Thus, to enable comparison, we will set $\theta = 0$ and therefore require $h_{\xi^0} > 0$; $h_{\xi^0} = 0$ is the point at which the memory is forgotten.

4 Stochastically Updating Synapses

The simplest, memory-related model of synaptic plasticity to consider is one in which synapses stochastically express synaptic plasticity, with probability p , in response to a plasticity induction signal (Tsodyks, 1990). Such a synapse does not possess any internal, plasticity-related states, and so the analysis in this case is greatly simplified. The single synapse transition matrices are simply

$$\mathbb{M}_- = \begin{pmatrix} 1 & p \\ 0 & 1-p \end{pmatrix}, \quad \mathbb{M}_+ = \begin{pmatrix} 1-p & 0 \\ p & 1 \end{pmatrix}, \quad \mathbb{M} = \begin{pmatrix} 1-\frac{1}{2}p & \frac{1}{2}p \\ \frac{1}{2}p & 1-\frac{1}{2}p \end{pmatrix}. \quad (4.1)$$

Since

$$\mathbb{M} = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} + \frac{1}{2}(1-p) \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \quad (4.2)$$

we have that

$$\mathbb{M}^m \rightarrow \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \quad (4.3)$$

as $m \rightarrow \infty$ and so $\mathbf{S}_i(t) \rightarrow \mathbf{A} \equiv \frac{1}{2}(1, 1)^\top$ as $t \rightarrow \infty$. In equilibrium, then, $\text{Prob}[S_i(t) = \pm 1] = 1/2$ so that both strength states are equiprobable. It is against this background equilibrium distribution of synaptic strengths given by $\mathbf{A} \otimes \cdots \otimes \mathbf{A}$ that the definite initial memory $\boldsymbol{\xi}^0$ is stored. Now, if synapse i receives a $\xi_i^0 = +1$ signal, then the distribution of its strengths immediately after storage is $\mathbb{M}_+ \mathbf{A} = \frac{1}{2}(1-p, 1+p)^\top$, while for $\xi_i^0 = -1$, this distribution is $\mathbb{M}_- \mathbf{A} = \frac{1}{2}(1+p, 1-p)^\top$. A synapse with $\xi_i^0 = +1$ contributes to $h(t)$ positively, while a synapse with $\xi_i^0 = -1$ contributes negatively. A weak synapse with $\xi_i^0 = -1$ therefore contributes positively to $h(t)$ while a strong synapse with $\xi_i^0 = -1$ contributes negatively, i.e. for $\xi_i^0 = -1$, the roles of strong and weak

strengths states are reversed in terms of their contributions to $h(t)$. Defining $\tilde{S}_i(t) = \xi_i^0 S_i(t)$, with $h(t) = \frac{1}{N} \sum_{i=1}^N \tilde{S}_i(t)$, and defining $\tilde{\mathbf{S}}_i(t) = \xi_i^0 \mathbf{S}_i(t)$, the initial distribution of $\tilde{\mathbf{S}}_i$ for $\xi_i^0 = +1$ immediately after the storage of $\boldsymbol{\xi}^0$ is $\frac{1}{2}(1 - p, 1 + p)^T$, and the initial distribution of $\tilde{\mathbf{S}}_i$ for $\xi_i^0 = -1$ immediately after storage is also $\frac{1}{2}(1 - p, 1 + p)^T$. These distributions are identical, so $h(t)$ is the sum over N independent, identically-distributed random variables immediately after the storage of $\boldsymbol{\xi}^0$. Moreover, each synapse is subsequently exposed to the same transition matrix \mathbb{M} at later memory storage steps. Thus, $h(t)$ remains for all time the sum over N independent, identically-distributed random variables. In order to avoid the repeated phrase, “immediately after the storage of $\boldsymbol{\xi}^0$ ”, we shall consider that $\boldsymbol{\xi}^0$ is stored at time $t = 0^-$ s, so that $t = 0$ s is the time immediately after storage and we can just refer to the state at $t = 0$ s. In the following, $h(t)$ refers to $h_{\boldsymbol{\xi}^0}$ with memories stored as a continuous-time process, while h_n will be understood to refer to $h_{\boldsymbol{\xi}^0}(t = n \text{ s})$ with memories stored at discrete, unit time steps. We will, however, often use h_0 and $h(0)$ interchangeably, for simplicity of notation.

5 Exact Transition Probabilities for h

We may use the observation that $h(t)$ is the sum over N independent, identically-distributed random variables $\tilde{S}_i(t)$ to calculate directly the mean and variance of the memory signal for an SNR determination of memory lifetimes (see Elliott and Lagogiannis, 2012). However, we proceed somewhat more gen-

erally using an approach that allows us also to obtain the MFPT determination of memory lifetimes. We do this by working directly with the transitions in the values of $h(t)$, or h_n , rather than the transitions in individual synaptic strengths.

If $h(t) = 2j/N - 1$, then there are precisely j of the $\tilde{S}(t)$ variables taking values $+1$ and $N - j$ taking values -1 . For the initial state at $t = 0^-$ s, then, we immediately have

$$\text{Prob} \left[h_0 = \frac{2j}{N} - 1 \right] = {}^N C_j \left(\frac{1+p}{2} \right)^j \left(\frac{1-p}{2} \right)^{N-j}, \quad (5.1)$$

with mean $\mathbb{E}[h_0] = p$ and variance $\text{Var}[h_0] = (1-p^2)/N$. If $h_n = 2j/N - 1$ and $h_{n+1} = 2i/N - 1$, then by considering the possible transitions of the j synapses with $\tilde{S} = +1$ and the $N - j$ synapses with $\tilde{S} = -1$, we see that the transition probability from $h_n = 2j/N - 1$ to $h_{n+1} = 2i/N - 1$ is

$$\begin{aligned} & \text{Prob} \left[h_{n+1} = \frac{2i}{N} - 1 \mid h_n = \frac{2j}{N} - 1 \right] \\ &= \sum_{k=0}^j {}^j C_k \left(\frac{p}{2} \right)^k \left(1 - \frac{p}{2} \right)^{j-k} {}^{N-j} C_{i-j+k} \left(\frac{p}{2} \right)^{i-j+k} \left(1 - \frac{p}{2} \right)^{N-i-k}, \end{aligned} \quad (5.2)$$

where the standard conventions regarding the binomial coefficients apply, i.e. ${}^a C_b = 0$ for $b < 0$ and $b > a$. Eq. (5.2) gives the one-memory-encoding-step transition matrix elements for h . We denote this matrix by \mathbb{T} . It is homogeneous, not depending on the encoding step n .

In Appendix B, using generating functions we compute \mathbb{T}^m , giving the m -memory-encoding-step transition matrix. From \mathbb{T}^m , we can obtain the conditional and unconditional means and variances in the memory signal after m

subsequent memories have been stored. Defining $q = 1 - p$, we obtain

$$\mathbb{E}[h_m|h_0] = q^m h_0, \quad (5.3)$$

$$\text{Var}[h_m|h_0] = \frac{1 - q^{2m}}{N}, \quad (5.4)$$

and

$$\mathbb{E}[h_m] = q^m p, \quad (5.5)$$

$$\text{Var}[h_m] = \frac{1 - q^{2m} p^2}{N} = \frac{1 - \mathbb{E}[h_m]^2}{N}. \quad (5.6)$$

The form for $\text{Var}[h_m]$ indicates the absence of any covariance terms between pairs of synapses' strengths. In continuous time, the equivalent results are

$$\mathbb{E}[h(t)|h(0)] = e^{-(1-q)rt} h(0), \quad (5.7)$$

$$\text{Var}[h(t)|h(0)] = \frac{1 - e^{-(1-q^2)rt}}{N} + \left[e^{-(1-q^2)rt} - e^{-2(1-q)rt} \right] h(0)^2, \quad (5.8)$$

and

$$\mathbb{E}[h(t)] = e^{-(1-q)rt} p, \quad (5.9)$$

$$\begin{aligned} \text{Var}[h(t)] &= \frac{1 - e^{-2(1-q)rt} p^2}{N} + \left(1 - \frac{1}{N}\right) p^2 \left[e^{-(1-q^2)rt} - e^{-2(1-q)rt} \right] \\ &= \frac{1 - \mathbb{E}[h(t)]^2}{N} + \left(1 - \frac{1}{N}\right) p^2 \left[e^{-(1-q^2)rt} - e^{-2(1-q)rt} \right]. \end{aligned} \quad (5.10)$$

The second term in the expression for $\text{Var}[h(t)]$ reflects the continuous-time-induced covariance between changes in pairs of synapses' strengths, despite their independent updates [see Elliott and Lagogiannis (2012) for a full discussion of this phenomenon]. We may use these discrete- and continuous-time results to obtain SNR memory lifetimes. The presence of a covariance term in

$\text{Var}[h(t)]$ and its absence in $\text{Var}[h_m]$ lead to different SNR lifetimes depending on whether we use a continuous-time approach or a discrete-time approach. For longer lifetimes, the difference is usually negligible, but for shorter lifetimes, it can be significant.

In the underlying tensor product space, the one-memory-encoding-step operator is $\mathbb{M}_N = \mathbb{M} \otimes \cdots \otimes \mathbb{M}$, while in terms of the transitions in the neuronal output h itself, the one-memory-encoding-step matrix is \mathbb{T} . We may therefore use \mathbb{T} to rewrite τ_{mfpt} in Eq. (3.7) or Eq. (3.9), removing the tensor products. The equivalent projection operator onto states with $h > 0$ is simply the $(N + 1) \times (N + 1)$ diagonal matrix,

$$\mathbb{P}_{h>0} = \text{diag}\left\{ \underbrace{0, \dots, 0}_{h \leq 0 \text{ states}} \mid \underbrace{1, \dots, 1}_{h > 0 \text{ states}} \right\}, \quad (5.11)$$

where we consider a canonical ordering of entries for $h = 2i/N - 1$ for $i = 0, \dots, N$. There are either $(N + 1)/2$ or $N/2$ entries of unity on the diagonal of $\mathbb{P}_{h>0}$ depending on whether N is odd or even, respectively. Writing $\mathbf{P}(h_0)$ as a vector denoting a particular, definite initial value for h_0 , so with an entry of unity at position $i = \frac{N}{2}(1 + h_0)$ and entries of zero elsewhere, and writing $\boldsymbol{\nu}$ as an $(N + 1)$ -dimensional vector with entries of unity everywhere, Eq. (3.7) translates directly into

$$\tau_{\text{mfpt}}(h_0) = \boldsymbol{\nu}^T [\mathbb{I}_{N+1} - \mathbb{P}_{h>0} \mathbb{T} \mathbb{P}_{h>0}]^{-1} [\mathbb{P}_{h>0} \mathbf{P}(h_0)], \quad (5.12)$$

and of course, the continuous-time form in Eq. (3.9) is identical but with a rate factor present. Even for sizeable values of N , up to around 10,000, the inverse matrix in Eq. (5.12) may be computed numerically. With the form $\mathbb{P}_{h>0} \mathbb{T} \mathbb{P}_{h>0}$

rather than $\mathbb{P}_{h>0}\mathbb{T}$, the dimensionality of the problem is reduced, for odd N for example, from $(N + 1) \times (N + 1)$ to $(N + 1)/2 \times (N + 1)/2$.

We may obtain an alternative expression for $\tau_{\text{mfpt}}(h_0)$. Writing $\mathbb{T}_{++} = \mathbb{P}_{h>0}\mathbb{T}\mathbb{P}_{h>0}$ as a convenient short-hand, we have that

$$\begin{aligned} \tau_{\text{mfpt}}(h_0) &= \boldsymbol{\nu}^T [\mathbb{I}_{N+1} + \mathbb{T}_{++} + (\mathbb{T}_{++})^2 + \dots] [\mathbb{P}_{h>0}\mathbf{P}(h_0)] \\ &= \boldsymbol{\nu}^T \mathbb{P}_{h>0}\mathbf{P}(h_0) \\ &\quad + \boldsymbol{\nu}^T [\mathbb{I}_{N+1} + \mathbb{T}_{++} + (\mathbb{T}_{++})^2 + \dots] \mathbb{T}_{++} [\mathbb{P}_{h>0}\mathbf{P}(h_0)]. \end{aligned} \tag{5.13}$$

The vector $\mathbb{T}_{++} [\mathbb{P}_{h>0}\mathbf{P}(h_0)]$ is the distribution of states of h after one additional memory has been stored, with $h \leq 0$ states projected out. Hence, assuming that $h_0 > 0$, we may write

$$\tau_{\text{mfpt}}(h_0) - 1 = \sum_{h'>0} \tau_{\text{mfpt}}(h') \text{Prob}[h'|h_0] \tag{5.14}$$

as an alternative form, with $\text{Prob}[h'|h_0]$ given by Eq. (5.2). A generalisation of this form will be extremely useful later.

6 Continuum Limit: Memory Nearly on a Spring

If N is large enough, then we may consider a continuum limit for the neuronal output $h(t)$ or h_m . This consists merely of replacing the binomial distribution by the Gaussian distribution for N large enough. In this limit, we must have

a Gaussian distribution with mean and variance given by Eqs. (5.5) and (5.6) for $m = 0$ for the initial distribution of h_0 , so

$$\text{Prob}[h_0] = \sqrt{\frac{N}{2\pi(1-p^2)}} \exp\left[-\frac{N(h_0-p)^2}{2(1-p^2)}\right], \quad (6.1)$$

and a conditional Gaussian distribution with conditional mean and variance given by Eqs. (5.3) and (5.4) for $m = 1$ for the transition probabilities, so

$$\text{Prob}[h_{n+1}|h_n] = \sqrt{\frac{N}{2\pi(1-q^2)}} \exp\left[-\frac{N(h_{n+1}-qh_n)^2}{2(1-q^2)}\right]. \quad (6.2)$$

We may directly confirm that this limit respects the Markovian property, namely

$$\int dh_{n+1} \text{Prob}[h_{n+2}|h_{n+1}] \text{Prob}[h_{n+1}|h_n] = \text{Prob}[h_{n+2}|h_n], \quad (6.3)$$

and so in general,

$$\text{Prob}[h_{n+m}|h_n] = \sqrt{\frac{N}{2\pi(1-q^{2m})}} \exp\left[-\frac{N(h_{n+m}-q^m h_n)^2}{2(1-q^{2m})}\right], \quad (6.4)$$

with

$$\text{Prob}[h_m] = \sqrt{\frac{N}{2\pi(1-q^{2m}p^2)}} \exp\left[-\frac{N(h_m-q^m p)^2}{2(1-q^{2m}p^2)}\right]. \quad (6.5)$$

We may in principle compute $\text{Prob}[h(t)|h(0)]$ and $\text{Prob}[h(t)]$ via Poisson summation over m (with $n = 0$) in Eqs. (6.4) and (6.5), but the sums cannot be evaluated in closed form. As we already know $\mathbb{E}[h(t)|h(0)]$ and $\text{Var}[h(t)|h(0)]$ from Eqs. (5.7) and (5.8), it is extremely tempting to write

$$\text{Prob}[h(t)|h(0)] \stackrel{?}{=} \sqrt{\frac{1}{2\pi \text{Var}[h(t)|h(0)]}} \exp\left\{-\frac{[h(t) - \mathbb{E}[h(t)|h(0)]]^2}{2 \text{Var}[h(t)|h(0)]}\right\}. \quad (6.6)$$

However, this cannot be correct since the only stationary, Markovian, Gaussian process is (essentially) the Ornstein-Uhlenbeck (OU) process (see, for example, van Kampen, 1992) and if Eq. (6.6) were correct, then we would have another distinct such process. We may, though, consider the limit of small p in Eqs. (5.7) and (5.8) so that we can write $1 - q^2 = 1 - (1 - p)^2 \approx 2p$. Then

$$\begin{aligned}\mathbb{E}[h(t)|h(0)] &= e^{-prt} h(0), \\ \text{Var}[h(t)|h(0)] &\approx \frac{1 - e^{-2prt}}{N},\end{aligned}\tag{6.7}$$

and we recognise these as precisely the conditional mean and variance of an OU process, so that we may write

$$\text{Prob}[h(t)|h(0)] \approx \sqrt{\frac{N}{2\pi(1 - e^{-2prt})}} \exp\left\{-\frac{N[h(t) - e^{-prt}h(0)]^2}{2(1 - e^{-2prt})}\right\}.\tag{6.8}$$

with

$$\text{Prob}[h(t)] \approx \sqrt{\frac{N}{2\pi(1 - e^{-2prt}p^2)}} \exp\left\{-\frac{N[h(t) - e^{-prt}p]^2}{2(1 - e^{-2prt}p^2)}\right\}.\tag{6.9}$$

The OU process may be thought of as describing the dynamics of a spring subject to random, Gaussian impulses. Thus, in this limit with $p \ll 1$ and $N \gg 1$, the initial memory encoding event at $t = 0^-$ s extends the spring to a length, on average, of p units, and the restoring force returns the spring back to zero length but subject to random fluctuations, giving an asymptotic variance of $1/N$.

The transition probability for this OU process satisfies the backward Kolmogorov or Fokker-Planck equation,

$$\left[\frac{\partial}{\partial(rt)} + ph(0)\frac{\partial}{\partial h(0)} - \frac{p}{N}\frac{\partial^2}{\partial h(0)^2}\right]\text{Prob}[h(t)|h(0)] = 0.\tag{6.10}$$

From this, we may obtain the MFPT for the system to escape from an interval for the first time by erecting absorbing boundaries. Erecting a boundary at $h = \theta \equiv 0$ and a second boundary at $h = \hat{\theta} > 0$ and then taking the limit $\hat{\theta} \rightarrow \infty$ so that it becomes irrelevant, we obtain the MFPT for a system initially with $h(0) = x > 0$ to be absorbed at $h = 0$. The calculation is standard and the details of similar calculations may be found elsewhere (Elliott, 2011a). The MFPT satisfies the differential equation

$$-\frac{1}{r} = -px\tau'_{\text{mfpt}}(x) + \frac{p}{N}\tau''_{\text{mfpt}}(x), \quad (6.11)$$

subject to the boundary conditions $\tau_{\text{mfpt}}(0) = 0$ and $\tau_{\text{mfpt}}(\hat{\theta}) = 0$. The solution, with $\hat{\theta} \rightarrow \infty$, is

$$r\tau_{\text{mfpt}}(x) = \frac{1}{2p} \left\{ \pi \operatorname{erfi} \left(\sqrt{\frac{N}{2}}x \right) - Nx^2 {}_2F_2 \left(\{1, 1\}; \left\{ \frac{3}{2}, 2 \right\}, \frac{Nx^2}{2} \right) \right\}, \quad (6.12)$$

where erfi is the imaginary error function and ${}_2F_2$ is a hypergeometric function.

If the firing threshold θ is non-zero, then this solution simply becomes

$$\tau_{\text{mfpt}}(x; \theta) = \tau_{\text{mfpt}}(x) - \tau_{\text{mfpt}}(\theta). \quad (6.13)$$

For large Nx^2 , there is an extremely delicate cancellation between the two terms on the right hand side of Eq. (6.12), leaving a logarithmic behaviour.

We find that

$$r\tau_{\text{mfpt}}(x) \sim \frac{1}{2p} \left[\log_e \frac{Nx^2}{2} - \psi\left(\frac{1}{2}\right) \right], \quad (6.14)$$

for large Nx^2 , where ψ is the digamma function, with $\psi(\frac{1}{2}) = -\hat{\gamma} - 2 \log_e 2$ where $\hat{\gamma} \approx 0.5772$ is Euler's constant. With non-zero θ , this asymptotic be-

haviour is just

$$r\tau_{\text{mfpt}}(x) \sim \frac{1}{2p} \left[\log_e \frac{Nx^2}{2} - \psi\left(\frac{1}{2}\right) - 2pr\tau_{\text{mfpt}}(\theta) \right] \quad (6.15)$$

$$\sim \frac{1}{p} \log_e \left(\frac{x}{\theta} \right), \quad (6.16)$$

where the second line follows in the case that $N\theta^2$ is large enough. Notice that the N -dependence has entirely disappeared in Eq. (6.16); we shall return to this issue later.

Taking $\text{Var}[h(t)|h(0)] \approx 1/N$ and writing $\tau_{\text{snr}}(x)$ for the continuous-time SNR determination of memory lifetimes with $h(0) = x$, we find that

$$r\tau_{\text{snr}}(x) \approx \frac{1}{2p} \left[\log_e \frac{Nx^2}{2} + \log_e 2 \right]. \quad (6.17)$$

If instead we require that the SNR falls to some general value ϑ rather than $\vartheta = 1$ specifically, then this result is modified, becoming

$$r\tau_{\text{snr}}(x) \approx \frac{1}{2p} \left[\log_e \frac{Nx^2}{2} + \log_e 2 - 2 \log_e \vartheta \right]. \quad (6.18)$$

The results in Eqs. (6.14) and (6.17) are identical up to an additive constant, although their respective ranges of validity may differ, depending on parameter values. We shall compare $\tau_{\text{mfpt}}(x)$ and $\tau_{\text{snr}}(x)$ both qualitatively and quantitatively when we are in possession of an expression for $\tau_{\text{mfpt}}(x)$ that does not depend on the small p limit.

Although we have taken the Gaussian limit of the state and transition probabilities in Eqs. (6.1) and (6.2) to define a continuum limit for h for $N \gg 1$, we note that these expressions are also extremely good approximations to the exact, binomial state and transition probabilities in Eqs. (5.1) and (5.2) for

discrete values of h . The conditions for this approximation are usually stated as requiring $N \gtrsim 20$ and both $Np \gtrsim 5$ and $Nq \gtrsim 5$. For numerical efficiency, we will frequently use this approximation for discrete h calculations when it is valid.

7 Mean First Passage Time for Continuum h

The expression for the MFPT in Eq. (5.14) for discrete h generalises immediately to a continuum h version. In continuous time, τ_{mfpt} becomes $r\tau_{\text{mfpt}}$, so for some definite, initial state $h(0) = x > 0$, we have

$$\tau_{\text{mfpt}}(x) - \frac{1}{r} = \int_0^\infty dy \tau_{\text{mfpt}}(y) \text{Prob}[y|x], \quad (7.1)$$

which is an integral equation with kernel $\text{Prob}[y|x]$ given by the Gaussian form in Eq. (6.2). We will write $K_0(y|x) = \text{Prob}[y|x]$ for simplicity, i.e.

$$K_0(y|x) = \sqrt{\frac{N}{2\pi(1-q^2)}} \exp\left[-\frac{N(y-qx)^2}{2(1-q^2)}\right]. \quad (7.2)$$

We now set $r = 1$, without loss of generality, so that we have 1 rather than $1/r$ on the left hand side of Eq. (7.1). The form in Eq. (7.1) is the standard expression for MFPTs for processes with jumps (van Kampen, 1992). That jump processes are present is clear from the discrete h process. For N synapses each with strength-change probability $\frac{1}{2}p$ upon memory storage, on average we expect $\frac{1}{2}Np$ synapses to change strength per memory encoding step. The neuronal output h may therefore be subject to large jumps and, in particular, it may jump over the $h = 0$ firing threshold. These jumps are automatically

accounted for in the discrete h form in Eq. (5.14) (and in the earlier form in Eq. (3.9)), of which the integral equation in Eq. (7.1) is the continuum limit.

Unfortunately, this integral equation is rarely soluble in closed form, although standard numerical methods are available. The Neumann series solution of Eq. (7.1) is

$$\tau_{\text{mfpt}}(x) = 1 + \int_0^\infty dy_1 K_0(y_1|x) + \int_0^\infty \int_0^\infty dy_2 dy_1 K_0(y_2|y_1) K_0(y_1|x) + \dots, \quad (7.3)$$

although typically the integrals cannot be evaluated nor the resulting series, if available, re-summed. However, we may make considerable progress by replacing the particular Gaussian kernel $K_0(y|x)$ by a simpler one that has identical first- and second-order conditional moments. For N large enough, $K_0(y|x)$ is a very narrow, nearly Dirac-delta-function-like spike. In this limit, its precise shape should not be too important, and provided that we use a replacement kernel with identical lower-order conditional moments, we would expect to obtain very good approximations to the exact solution for $\tau_{\text{mfpt}}(x)$ with the original Gaussian kernel. The trick is to pick a replacement kernel that allows us to recast the integral equation in Eq. (7.1) as a differential equation that can be solved, by a power series method if necessary. We will consider two such replacements.

7.1 Formal Kernel Replacement

We first consider a replacement based on a formal expansion in the Dirac delta function, $\delta(x)$, and its derivatives. Writing

$$K_1(y|x) = \delta(y-x) + px\delta'(y-x) + \frac{1}{2} \left(p^2x^2 + \frac{1-q^2}{N} \right) \delta''(y-x), \quad (7.4)$$

it is a simple matter to check that $\int_{-\infty}^{+\infty} dy y^m K_1(y|x)$ gives the correct conditional moments for $m = 0, 1$ and 2 , of 1 , qx and $q^2x^2 + (1-q^2)/N$ (and hence conditional variance $(1-q^2)/N$), respectively. The advantage of this kernel is that when we insert it into Eq. (7.1), we automatically obtain a differential equation because of the Dirac delta functions. For $x < 0$, we obtain the solution $\tau_{\text{mfpt}}(x) = 1$, while for $x > 0$, we have

$$-1 = -px\tau'_{\text{mfpt}}(x) + \frac{1}{2} \left(p^2x^2 + \frac{1-q^2}{N} \right) \tau''_{\text{mfpt}}(x). \quad (7.5)$$

For small p , the p^2x^2 term may be dropped and $1-q^2 \approx 2p$ and so we obtain Eq. (6.11). The solution in this limit is then

$$\tau_{\text{mfpt}}(x) = A + B \sqrt{\frac{\pi}{2N}} \operatorname{erfi} \left(\sqrt{\frac{N}{2}} x \right) - \frac{Nx^2}{2p} {}_2F_2 \left(\{1, 1\}; \left\{ \frac{3}{2}, 2 \right\}, \frac{Nx^2}{2} \right). \quad (7.6)$$

Usually, the constants A and B are uniquely determined by substituting this solution back into the integral equation for the MFPT in Eq. (7.1) with kernel $K_1(y|x)$. However, because of the purely formal character of this kernel, we merely obtain trivial identities that do not fix these constants. We may instead determine A and B by requiring continuity at $x = 0$, thus giving $A = 1$, and by observing that for large x , $\tau_{\text{mfpt}}(x) \sim \frac{1}{p} \log_e x$ is a solution of Eq. (7.5) for small p . Thus, we may set B so that this logarithmic behaviour emerges from

the above-mentioned delicate cancellation between erfi and ${}_2F_2$, so we require $B = \frac{1}{p} \sqrt{\frac{N\pi}{2}}$. Hence, we obtain, for small p with kernel $K_1(y|x)$,

$$\tau_{\text{mfpt}}(x) = 1 + \frac{1}{2p} \left\{ \pi \operatorname{erfi} \left(\sqrt{\frac{N}{2}} x \right) - Nx^2 {}_2F_2 \left(\{1, 1\}; \left\{ \frac{3}{2}, 2 \right\}, \frac{Nx^2}{2} \right) \right\}, \quad (7.7)$$

which is identical to Eq. (6.12) (with $r = 1$) up to an additive constant.

By direct substitution we may verify that the full solution of Eq. (7.5) is the rather unwieldy

$$\begin{aligned} \tau_{\text{mfpt}}(x) = & A + Bx \left(1 + \frac{Np^2x^2}{1-q^2} \right)^{1+\frac{1}{p}} {}_2F_1 \left(\left\{ 1, \frac{3}{2} + \frac{1}{p} \right\}, \frac{3}{2}, -\frac{Np^2x^2}{1-q^2} \right) \\ & - \frac{Nx^2}{1-q^2} {}_3F_2 \left(\left\{ 1, 1, \frac{1}{2} - \frac{1}{p} \right\}; \left\{ \frac{3}{2}, 2 \right\}, -\frac{Np^2x^2}{1-q^2} \right). \end{aligned} \quad (7.8)$$

By comparing the power series in the small p limit, we find a direct correspondence between the particular solutions and the non-constant general solutions in Eqs. (7.6) and (7.8). Again, continuity fixes $A = 1$. For large x , the logarithmic behaviour of the solution of Eq. (7.5) becomes $\tau_{\text{mfpt}}(x) \sim \frac{2}{2p+p^2} \log_e x$. The constant B is again set by imposing this logarithmic behaviour in Eq. (7.8) for large x , resulting in the equally unwieldy

$$B = \frac{2^{\frac{2}{p}-1} \sqrt{N} \pi^2 \sec^2 \frac{\pi}{p}}{\sqrt{1-q^2} \Gamma\left(\frac{1}{2} - \frac{1}{p}\right)^2 \Gamma\left(\frac{2}{p}\right)} \quad (7.9)$$

$$\sim \sqrt{\frac{N\pi}{2}} \frac{1}{p} \left(1 + \frac{1}{8}p + \frac{9}{128}p^2 + \dots \right), \quad (7.10)$$

where $\Gamma(x)$ is the gamma function. The asymptotic behaviour of the full solution of Eq. (7.5), including constants, is then found to be

$$\tau_{\text{mfpt}}(x) \sim 1 + \frac{1}{2p+p^2} \left[\log_e \frac{Np^2x^2}{1-q^2} - \psi\left(\frac{1}{2}\right) + \psi\left(-\frac{1}{2} - \frac{1}{p}\right) \right] \quad (7.11)$$

$$\xrightarrow{\text{small } p} \frac{1}{2p} \left[\log_e \frac{Nx^2}{2} - \psi\left(\frac{1}{2}\right) \right], \quad (7.12)$$

with the small p limit agreeing with Eq. (6.14).

7.2 Laplace Kernel Replacement

A second kernel that permits the derivation of a differential equation for $\tau_{\text{mfpt}}(x)$ is the Laplacian rather than Gaussian form,

$$K_2(y|x) = \frac{\gamma}{2} e^{-\gamma|y-qx|}, \quad (7.13)$$

where $\gamma^2 = 2N/(1 - q^2)$. Again, the conditional mean and variance are the same as those for $K_0(y|x)$. This kernel satisfies the differential equation

$$\left(\frac{\partial^2}{\partial x^2} - q^2 \gamma^2 \right) K_2(y|x) = -q^2 \gamma^2 \delta(qx - y), \quad (7.14)$$

and hence we can employ the differential operator on the left hand side to kill the kernel under the integral sign in Eq. (7.1), obtaining the differential equation

$$\tau_{\text{mfpt}}''(x) = q^2 \gamma^2 [\tau_{\text{mfpt}}(x) - 1 - H(x) \tau_{\text{mfpt}}(qx)], \quad (7.15)$$

where $H(x)$ is the Heaviside step function. For $x < 0$, insisting on a finite solution as $x \rightarrow -\infty$, we have the solution

$$\tau_{\text{mfpt}}(x) = 1 + [\tau_{\text{mfpt}}(0) - 1] e^{+q\gamma x} \quad (x < 0), \quad (7.16)$$

where $\tau_{\text{mfpt}}(0)$ is to be determined. For large $x > 0$, if we assume that $\tau_{\text{mfpt}}''(x) \rightarrow 0$, then the asymptotic solution must satisfy

$$\tau_{\text{mfpt}}(x) - \tau_{\text{mfpt}}(qx) = 1, \quad (7.17)$$

which has solution $\tau_{\text{mfpt}}(x) = C - (\log_e x)/(\log_e q)$ with constant C . Since such a solution indeed satisfies $\tau_{\text{mfpt}}''(x) \rightarrow 0$ as $x \rightarrow \infty$, this asymptotic behaviour

is consistent. As $-\log_e q \approx p$ for small p , we recognise the limiting behaviour of the OU solution. For $x > 0$, the full solution of Eq. (7.15) must be obtained as a power series, with the result

$$\tau_{\text{mfpt}}(x) = A + B \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+1}}{(2m+1)!} (q; q^2)_m - \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+2}}{(2m+2)!} (q^2; q^2)_m, \quad (7.18)$$

where $(a; q)_k \equiv \prod_{l=0}^{k-1} (1 - a q^l)$ with $(a; q)_0 \equiv 1$ is the q -Pochhammer symbol. Both series are absolutely convergent for all x . By considering the small p limit of the powers series for the particular solution and the non-constant general solution, we again find the OU limit:

$$\frac{1}{q\gamma} \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+1}}{(2m+1)!} (q; q^2)_m \sim \sqrt{\frac{\pi}{2N}} \operatorname{erfi} \left(\sqrt{\frac{N}{2}} x \right), \quad (7.19)$$

$$- \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+2}}{(2m+2)!} (q^2; q^2)_m \sim -\frac{Nx^2}{2p} {}_2F_2 \left(\left\{ 1, 1 \right\}; \left\{ \frac{3}{2}, 2 \right\}, \frac{Nx^2}{2} \right), \quad (7.20)$$

at leading order in p . (We do not pull a factor of $(q\gamma)^{-1}$ out of the constant B in Eq. (7.18) in order to make results below somewhat cleaner.) The constants A and B are determined by plugging this solution back into the original integral equation with the $K_2(y|x)$ kernel. Writing

$$\alpha_i = q^i \sum_{m=0}^{\infty} q^{2m} (q^i; q^2)_m, \quad (7.21)$$

we find that $B = (1 - \alpha_2)/(1 - \alpha_1)$ and $A = 1 + B$. Using the theory of q -series, we may explicitly evaluate the sum defining α_i . The details may be found in Appendix C. The result is

$$\alpha_i = 1 - (q^i; q^2)_{\infty}. \quad (7.22)$$

The full solution is therefore

$$\begin{aligned} \tau_{\text{mfpt}}(x) = & 1 + \frac{(q^2; q^2)_\infty}{(q; q^2)_\infty} \left[1 + \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+1}}{(2m+1)!} (q; q^2)_m \right] \\ & - \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+2}}{(2m+2)!} (q^2; q^2)_m. \end{aligned} \quad (7.23)$$

The behaviours of $(q; q^2)_\infty$ and $(q^2; q^2)_\infty$ for $q \sim 1$, i.e. $p \sim 0$, are well-known.

For $q = e^{-\rho}$, and thus $\rho = -\log_e(1-p)$, we have (Watson, 1936)

$$(q; q^2)_\infty \sim \sqrt{2} \exp\left(-\frac{\pi^2}{12\rho} - \frac{\rho}{24}\right) + o(1), \quad (7.24)$$

$$(q^2; q^2)_\infty \sim \sqrt{\frac{\pi}{\rho}} \exp\left(-\frac{\pi^2}{12\rho} + \frac{\rho}{12}\right) + o(1). \quad (7.25)$$

Individually, neither expression has a power series expansion around $\rho = 0$ and therefore around $p = 0$, but their ratio does. We have

$$\frac{(q^2; q^2)_\infty}{(q; q^2)_\infty} = \sqrt{\frac{\pi}{2p}} \left(1 - \frac{1}{8}p - \frac{13}{384}p^2 - \frac{49}{3072}p^3 - \dots \right). \quad (7.26)$$

Because the particular solution and the non-constant general solution in Eq. (7.23) do not appear to be standard special functions (they are neither hypergeometric nor q -hypergeometric functions), their asymptotic behaviours for large x are unknown, although we do know that there is a cancellation, leaving the logarithmic behaviour $\tau_{\text{mfpt}}(x) \sim C - (\log_e x)/(\log_e q)$. We also observe that the constant C is γ -dependent and thus N -dependent, but this dependence can be eliminated by instead writing $\tau_{\text{mfpt}}(x) \sim C' - (\log_e \gamma x)/(\log_e q)$ for some other constant C' ; indeed, only the combination γx appears in Eq. (7.23). We may at least obtain the asymptotic behaviour in principle order-by-order in p , by expanding the two power series in Eq. (7.23) in orders of p , obtaining increasingly higher order hypergeometric functions, whose asymptotic behaviours are

known. We give the expansion of Eq. (7.23) up to second order in Appendix D. We do not labour these tedious and unenlightening calculations here, but instead just give the asymptotic behaviour for the first few, calculated orders in p :

$$\begin{aligned} \tau_{\text{mfpt}}(x) \sim & \frac{3}{4} - \frac{5}{72}p + \sqrt{\frac{\pi}{2p}} \left(1 - \frac{1}{8}p\right) \\ & + \frac{1}{2p} \left(1 - \frac{1}{2}p - \frac{1}{12}p^2\right) \left[\log_e \frac{Nx^2}{2} - \psi\left(\frac{1}{2}\right)\right]. \end{aligned} \quad (7.27)$$

Again the small p limit is correct, giving Eq. (6.14).

For the Laplacian $K_2(y|x)$ kernel, we may also explicitly evaluate all the terms in the Neumann series in Eq. (7.3) and obtain a different representation of the solution in Eq. (7.23). We sketch the details in Appendix E, giving here only the final result,

$$\tau_{\text{mfpt}}(x) = 1 + \frac{(q^2; q^2)_\infty}{(q; q^2)_\infty} + (q^2; q^2)_\infty \sum_{n=0}^{\infty} \frac{1 - \exp(-\gamma q^{n+1}x)}{(q^2; q^2)_n}. \quad (7.28)$$

By expanding $\exp(-\gamma q^{n+1}x)$ as a power series in x and summing over n using the methods in Appendix C, we may show that the form in Eq. (7.28) is identical to that in Eq. (7.23). Although necessarily identical, these two representations have differing rates of convergence and each is therefore useful for numerical evaluation in differing regions of parameter space. For example, for p away from zero, q^n rapidly drops to zero, so only a few terms in Eq. (7.28) need be retained, even for large x ; in contrast, for large x , many terms in Eq. (7.23) must be evaluated. However, for p close to zero, q^n falls more slowly, so many terms in Eq. (7.28) must be computed and often Eq. (7.23) is numerically more efficient.

8 Results: Comparison Between Analytical, Numerical and Simulation Methods

Simulations are based on protocols described elsewhere (Elliott and Lagogiannis, 2012) but modified to extract MFPTs rather than SNRs. Extracting MFPTs requires running simulations until the memory signal has dropped to or below firing threshold. Critically, simulations must not be terminated after some fixed time step, as this would introduce a systematic bias that would lead to an underestimate of MFPTs. Computationally, extracting MFPTs from simulations is therefore more time-consuming than extracting SNRs. We average results over many simulations in order to obtain the *mean* first passage time rather than just a particular first passage time. Since each simulation constitutes a particular realisation of the randomly-generated, initial memory ξ^0 , the MFPT that emerges from simulation is not $\tau_{\text{mfpt}}(h_0)$ for some definite initial memory signal h_0 , but rather an average, $\langle \tau_{\text{mfpt}}(h_0) \rangle$, over all possible initial memories ξ^0 (or at least a very good sample thereof) leading to $h_0 > 0$. (If $h_0 \leq 0$, the simulation is of course immediately terminated and its particular first passage time set to zero.) We must therefore typically compare MFPTs from simulation to $\langle \tau_{\text{mfpt}}(h_0) \rangle$ where the average is taken over the initial distribution of h_0 . For N large enough, however, the distribution of h_0 is sharply focused around $h_0 = p$, and so we would expect that $\langle \tau_{\text{mfpt}}(h_0) \rangle \rightarrow \tau_{\text{mfpt}}(p)$ as $N \rightarrow \infty$.

Below we will compare MFPTs and SNRs, but SNRs are based on the

ratio of the mean memory signal to its standard deviation, and the initial mean memory signal is just p . However, for discrete values of h , p will not usually be a possible value of h_0 . Given a set of discrete values of $\tau_{\text{mfpt}}(h_0)$, we therefore construct an interpolating function to obtain $\tau_{\text{mfpt}}(p)$. Interpolation avoids introducing the noise into the inferred value of $\tau_{\text{mfpt}}(p)$ that would result were we instead to use the value of $\tau_{\text{mfpt}}(h_0)$ with h_0 closest to p , and such a procedure would also be extremely sensitive to whether N is even or odd, as described next.

MFPTs for discrete h states can depend quite sensitively on whether N is even or odd, since this determines whether or not $h = 0$ is a permitted discrete state. We may see this sensitivity by explicitly considering the trivial, $p = 1$ limit of MFPTs. In this limit, synapses always change strength in the presence of an induction signal of the appropriate sign. For $p = 1$, we have that $q = 0$, and for the continuum case, the Gaussian $K_0(y|x)$ kernel in Eq. (7.2) loses its x -dependence and hence Eq. (7.1) becomes $\tau_{\text{mfpt}}(x) - 1 = \text{const.}$ for some constant, for $x > 0$. Thus, we see that $\tau_{\text{mfpt}}(x) = 2$ for $x > 0$, and this does not depend on N . In the discrete case, however, for $p = 1$, the distribution of strengths at any subsequent encoding step m after the storage of ξ^0 is given by a symmetric binomial distribution (see Appendix B, specifically Eq. (B.5) for $p = 1$). Thus, the probability, call it \mathbf{p} , that $h > 0$ after any subsequent encoding step is just

$$\mathbf{p} = \begin{cases} \frac{1}{2} & \text{for } N \text{ odd} \\ \frac{1}{2} \left(1 - \frac{1}{2^N} {}^N C_{N/2}\right) & \text{for } N \text{ even} \end{cases}, \quad (8.1)$$

where the additional term for even N removes the disallowed $h = 0$ state that exists for even N . For this $p = 1$ case, the MFPT is therefore equivalent to the mean number of tosses for a (possibly biased) coin to come up with, say, heads (probability $1 - \mathfrak{p}$) for the first time. This is simply the sum

$$1 \mathfrak{p}^0(1 - \mathfrak{p}) + 2 \mathfrak{p}^1(1 - \mathfrak{p}) + 3 \mathfrak{p}^2(1 - \mathfrak{p}) + \dots = \frac{1}{1 - \mathfrak{p}}. \quad (8.2)$$

This result can also be seen directly from Eq. (5.14) since $\sum_{h' > 0} \text{Prob}[h'|h] \equiv \mathfrak{p}$. For odd N for discrete h , the MFPT is therefore always 2, agreeing with the continuum limit, while for even N , the MFPT is

$$\frac{2}{1 + \frac{1}{2^N} {}^N C_{N/2}} \xrightarrow{N \rightarrow \infty} 2. \quad (8.3)$$

The even N values are always lower than the adjacent, odd N values, but the values necessarily converge as $N \rightarrow \infty$. This is true for $p \neq 1$, too. Below, therefore, we shall always use even N values for consistency, to avoid spurious “wobbles” in graphs, and because it is better slightly to under-estimate MFPTs in general than to over-estimate them.

First we validate our discrete analytical results, for both the exact binomial transition probabilities and their Gaussian approximation, by comparing them to simulation results. In Fig 1, we plot $\langle \tau_{\text{mfpt}}(h_0) \rangle$ as a function of N for different choices of the update probability p . Exact binomial results are determined for $N \leq 2.5 \times 10^3$ while their Gaussian approximations are determined for $p^{-1} \lesssim N \leq 10^4$. For smaller N , there are tiny discrepancies between these exact and approximate MFPTs, but as expected they converge and become indistinguishable for larger values of N . Also shown in the same figure are

simulation results for the particular values $N = 10^2$, $N = 10^3$, $N = 5 \times 10^3$ and $N = 10^4$. The agreement between simulation results and analytical results is exact.

Because numerical matrix inversion becomes increasingly time-consuming as N grows, we are essentially limited to a maximum value of N of around 10^4 in Fig. 1 for discrete values of h . For larger N , we must move to the continuum h limit and solve the integral equation for MFPTs for the Gaussian kernel $K_0(y|x)$ using standard numerical methods. In Fig. 2, we compare the discrete h results in Fig. 1 with the continuum, Gaussian kernel results. The discrete h results are patched together essentially seamlessly from the exact binomial results for smaller N and the approximate Gaussian results for larger N . As expected, the discrete and continuum results differ somewhat for smaller N , but as N increases, the continuum limit becomes an increasingly good approximation. We also show additional simulation results for $N = 5 \times 10^4$ and $N = 10^5$, observing essentially perfect agreement between the continuum results and the simulation results for these larger values of N . We do not consider larger values of N because they would be biologically implausible, although the agreement would of course be maintained.

Although simulations generate results for $\langle \tau_{\text{mfpt}}(h_0) \rangle$ where we average over the initial signal of the memory stored at $t = 0^-$ s, computing this average from complicated analytical results for different continuum kernels can be time-consuming. If we can work instead with $\tau_{\text{mfpt}}(p)$, where $\langle h_0 \rangle \equiv p$, then such averaging can be eliminated. This would in some very weak sense




Figure 1. Comparison between MFPTs for exact binomial transition probabilities, approximate Gaussian transition probabilities and simulation results for discrete values of the perceptron output h . For each value of the update probability p as indicated, we plot $\langle \tau_{\text{mfpt}}(h_0) \rangle$ against N for binomial transition probabilities (solid line), Gaussian transition probabilities (dashed line) and simulation results (circles). Exact binomial results are obtained only for $N \leq 2.5 \times 10^3$ while approximate Gaussian results are obtained for $N \leq 10^4$, although the validity of the approximation sets a lower bound on Np . Here we use $N \gtrsim 1/p$. Simulation results are obtained by averaging over $10^8/N$ trials for each value of N and choice of p ; the larger N , the less averaging is required to obtain good estimates of $\langle \tau_{\text{mfpt}}(h_0) \rangle$ due to higher levels of self-averaging for more synapses.




Figure 2. Comparison between MFPTs for discrete values of h and the continuum limit of h . Analytical results for discrete values of h (solid lines) are taken from Fig. 1 by stitching together exact binomial and approximate Gaussian results as N increases. Continuum results (dashed lines) are obtained by the numerical solution of the integral equation for MFPTs with the Gaussian $K_0(y|x)$ kernel. Additional simulation results (circles) are shown for $N = 5 \times 10^4$ and $N = 10^5$.

be a mean-field approximation, although one warranted by our expectation that for large N , the distribution of h_0 becomes increasingly sharply focused around p , rather than just being a convenient ansatz to facilitate or simplify calculations. We compare $\tau_{\text{mfpt}}(p)$ and $\langle \tau_{\text{mfpt}}(h_0) \rangle$ by plotting their ratio as a function of N in Fig. 3. Indeed, as N increases, this ratio asymptotes to unity, although as expected, smaller values of p require larger values of N for this asymptotic behaviour to emerge. Below, we may therefore use $\tau_{\text{mfpt}}(p)$ rather than $\langle \tau_{\text{mfpt}}(h_0) \rangle$. This will, furthermore, facilitate comparison between MFPTs and SNRs, since SNRs are based on a definite, mean initial memory signal of $h_0 = p$.

The Gaussian kernel $K_0(x|y)$ does not lend itself to any analytical understanding of the dependence of $\tau_{\text{mfpt}}(p)$ on the parameters p and N because numerical methods are necessary to solve the integral equation governing MFPTs. If the formal kernel $K_1(x|y)$ and the Laplace kernel $K_2(x|y)$ provide good replacements with satisfactory approximations to $\tau_{\text{mfpt}}(p)$ for the Gaussian kernel, then we may gain some analytical insight. In Fig. 4, we therefore plot results for all three kernels, together with the derived asymptotic or logarithmic behaviours of $\tau_{\text{mfpt}}(p)$ for the two replacement kernels. The very close agreement between the Gauss and Laplace results is both striking and satisfying. There is also good qualitative agreement between the Gauss and formal results, with the formal results following perfectly the trends in the Gauss results as a function of N . All three kernels provide results that differ essentially only by a small, fixed constant for any value of N . The logarithmic, asymptotic




Figure 3. Convergence of $\tau_{mfpt}(p)$ and $\langle \tau_{mfpt}(h_0) \rangle$ as N increases. Solid lines show $\tau_{mfpt}(p)/\langle \tau_{mfpt}(h_0) \rangle$ for continuum h , while dashed lines show the same ratio for discrete h . One set of dashed lines (those that start at lower values of N) are those corresponding to exact binomial results, while the other set (those that start at higher values of N when the approximation becomes valid) are those corresponding to the approximate Gaussian results.

results for the formal and Laplace kernels are also seen to match perfectly the exact forms as N grows. We note, however, that were we to use the logarithmic results to estimate MFPTs, we would typically obtain a significant underestimate of MFPTs, particularly for smaller N . Indeed, we would wrongly deduce that there is a minimum value of N only above which memory storage is possible. However, MFPTs are positive for all displayed values of N and so no such lower bound exists. For $p = 0.01$, for example, we would wrongly deduce that the minimum number of synapses required for successful initial memory encoding is around 2000 to 3000. Since neurons typically have at most a few thousand synapses, such an badly wrong estimate may lead to radically false conclusions about the viability of models with biologically realistic numbers of synapses. In fact, for around 2000 to 3000 synapses for $p = 0.01$, memory lifetimes gauged by MFPTs are already distinctly non-zero and positive.

The memory-on-a-spring, OU limit is a formal limit for $p \rightarrow 0$, based essentially on a diffusion approximation of the full dynamics. The computation of the MFPT in this OU limit by using an absorbing boundary implies smooth, continuous transitions without jump processes, with all intermediate states being visited as the system evolves towards, and is eventually absorbed by, the boundary. In Figs. 5 and 6, we compare MFPTs from this OU limit with those obtained from the formal, $K_1(y|x)$ kernel and the Laplace, $K_2(y|x)$ kernel, respectively, and also compare their asymptotic, logarithmic behaviours. The formal kernel is an expansion in conditional moments up to second order involved Dirac delta functions of zero width. Jump processes are therefore

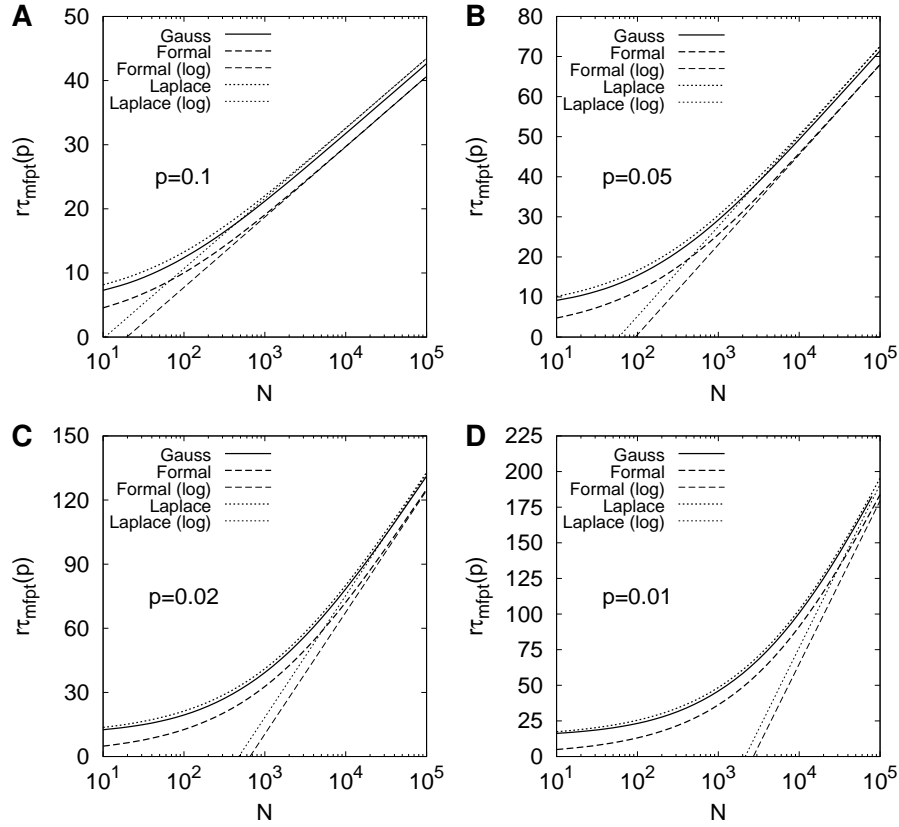


Figure 4. Comparison between continuum h MFPT results for the Gaussian $K_0(y|x)$ kernel, the formal $K_1(y|x)$ kernel and the Laplace $K_2(y|x)$ kernel, including the logarithmic forms for the formal and Laplace kernels. Gaussian kernel results are again obtained by numerical methods while formal and Laplace kernel results correspond to the analytical results discussed in the main text. The four panels show results for different update probabilities as indicated.

absent from this kernel and so we would expect increasingly good quantitative agreement between the formal kernel MFPTs and the OU results as p is reduced. Indeed we see explicitly this convergence in Fig. 5. The Laplace kernel, however, has a finite width and therefore retains jump processes and indeed, as we have seen in Fig. 4, provides an extremely good quantitative approximation to the Gaussian kernel. Although a small p limit will suppress jump processes, because the average number of strength changes per memory storage event, $\frac{1}{2}Np$, reduces as p reduces, increasing N will necessarily reinstate these jump processes. Thus, although we expect good qualitative agreement between the OU limit and the Laplace kernel, we should not expect that the MFPTs will converge for all N as p is reduced, except in the formal mathematical limit, $Np \rightarrow 0$. Fig. 6 bears out these expectations: we observe good qualitative agreement but with little if any hint of convergence of MFPTs for smaller p . The MFPTs in these two cases differ by an amount that is a roughly fixed constant, independent of N , for any given choice of p .

The asymptotic behaviour of MFPTs for the OU dynamics is particularly simple, given in Eq. (6.14). Its form is identical to memory lifetimes determined by SNRs in Eq. (6.17) up to an overall, p -dependent additive constant. Were we to plot $\tau_{\text{snr}}(p)$ on the graphs in Fig. 6, we would obtain a straight line of identical gradient to the asymptotic behaviour of $\tau_{\text{mfpt}}(p)$ for the OU process, but shifted overall downwards due to the additive constant. For the asymptotic OU process, we would deduce that the minimum number of synapses required for a non-zero memory lifetime is $N_{\text{min}}^{\text{ou}} = \exp(-\hat{\gamma})/2p^2$, where $\exp(-\hat{\gamma})/2 \approx$

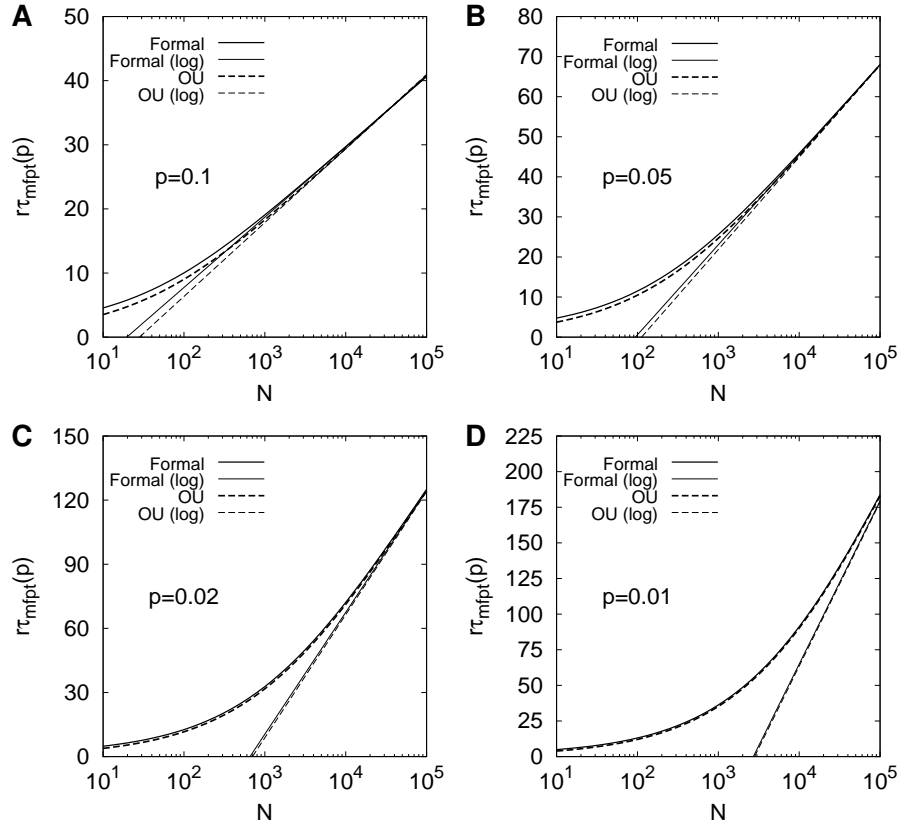


Figure 5. Comparison between the OU limit MFPT results and those for the formal $K_1(y|x)$ kernel, including their respective asymptotic, logarithmic behaviours.

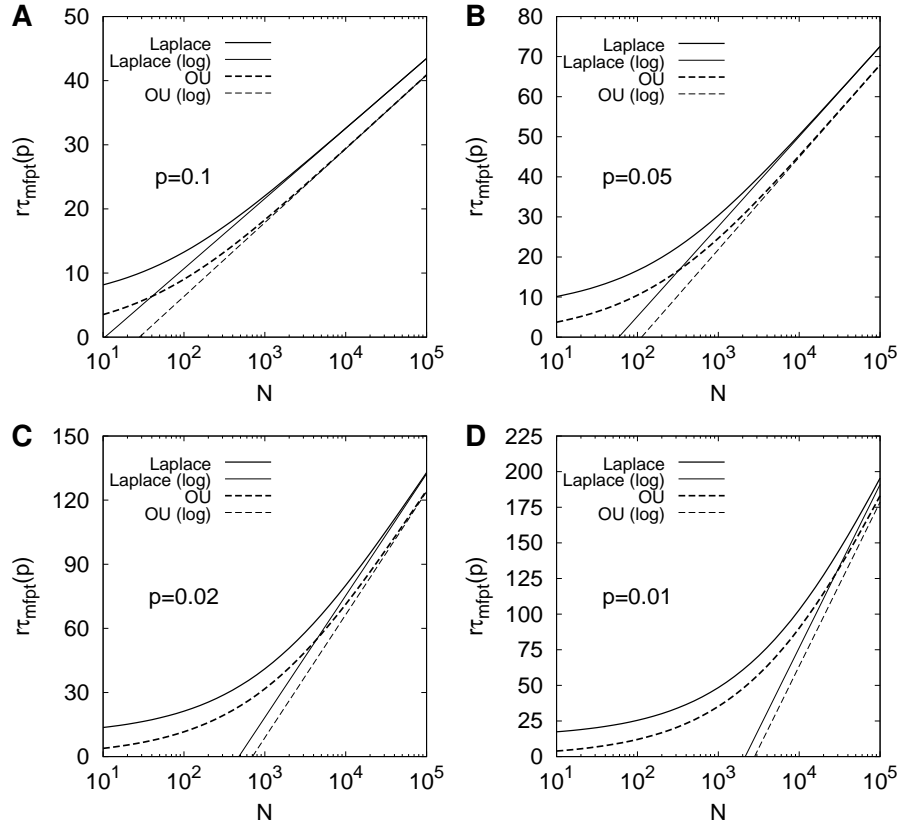


Figure 6. Comparison between the OU limit MFPT results and those for the Laplace $K_2(y|x)$ kernel, including their respective asymptotic, logarithmic behaviours.

0.28, or roughly $N_{\min}^{\text{ou}} \approx 1/4p^2$. We have seen above that this number is badly wrong, since the asymptotic behaviour is invalid for smaller values of N (for the OU limit but also for any kernel). For SNRs, we would deduce that this minimum number is $N_{\min}^{\text{snr}} = 1/p^2$, i.e. nearly four times larger than N_{\min}^{ou} . We point out that when computing SNRs, it is usually required that the initial SNR exceeds unity, i.e. $p\sqrt{N} > 1$, which leads precisely to $N > 1/p^2$. Of course, within an MFPT approach, such an initial requirement is not valid nor is it necessary, and this is reflected in the small N behaviour of MFPTs, and specifically in the fact that they are positive even for very small N . For MFPTs, all we require initially is that the neuron’s total input exceeds firing threshold. How much it exceeds threshold is irrelevant, for the Heaviside non-linearity will take care of the rest, setting the neuron’s actual output to unity. Arguments based on the size of the total input and of the resulting “vividness” of recall are therefore entirely fallacious in the presence of such a non-linearity.

Is it possible that this significant qualitative and quantitative disagreement between memory lifetimes gauged by MFPTs and SNRs owes its origin to an inappropriately chosen firing threshold, $\theta = 0$? An SNR criterion corresponds to the mean memory signal reaching zero at some definite number of standard derivations, typically at the one standard deviation level. Perhaps the one standard deviation level is not equivalent to a firing threshold $\theta = 0$, thus perhaps explaining the above disagreement between MFPT and SNR memory lifetimes? This is not, however, the case, as shown in Fig. 7. Indeed, the disagreement becomes far more profound, with far deeper and wide-ranging

implications, when $\theta > 0$. We see that the logarithmic, asymptotic behaviour is entirely destroyed when $\theta > 0$, and that $\tau_{\text{mfpt}}(p)$ asymptotes to the *constant* value $\frac{1}{p} \log_e \left(\frac{p}{\theta}\right)$ given by Eq. (6.16). Specifically, the N -dependence is essentially lost: increasing N will not increase this maximum possible memory lifetime. Although this result is derived from the OU limit in the presence of an absorbing boundary at $h = \theta > 0$, simulation results confirm that this behaviour is present away from the OU limit. A neuron with a non-zero firing threshold θ therefore has a fixed, maximum upper capacity to the memories that it can store, and increasing the number of its synapses will not increase this maximum. This asymptotic saturation of memory capacity occurs for *any* non-zero, positive value of θ , regardless of how small, because the asymptotic validity condition, that $N\theta^2$ is large enough, can always be achieved by taking N large enough. While the number of synapses required to achieve this asymptotic behaviour may not be biologically relevant, this upper bound on memory capacity nevertheless still exists. The logarithmic asymptotic behaviour at $\theta = 0$ is therefore entirely untypical and non-generic. Moreover, the $\theta \rightarrow 0$ limit is therefore discontinuous. If the overall $\log_e N$ SNR capacity of a neuron with binary synapses is usually regarded as a problem, then this upper bound on MFPT capacity with $\theta > 0$ is a catastrophe. We discuss this result, and the reasons for it, extensively in the next section.

However that may be, and despite the far deeper consequences, on a more parochial level, this asymptotic saturation shows that the difference in memory lifetimes between SNR and MFPT for $\theta = 0$ cannot be due to an inappro-




Figure 7. MFPTs for the OU limit in the presence of a possibly non-zero firing threshold θ . Solid lines show exact MFPTs for this OU limit for the values of θ as indicated. Dashed lines show the asymptotic forms of these MFPTs, derived in the large $N\theta^2$ limit. In this asymptotic limit, the N -dependence disappears entirely, so these dashed lines are constants at $\frac{1}{p} \log_e \left(\frac{p}{\theta} \right)$, to which the exact OU MFPTs asymptote as N grows for $\theta > 0$. Simulation results (circles) for $\langle \tau_{\text{mfpt}}(h_0) \rangle$ are also shown for comparison and validation of this asymptotic behaviour for $\theta = 0$, $\theta = p/5$ and $\theta = 2p/5$. For $N = 100$ there appear to be only two simulation datum points. However, the results for $\theta = 0$ and $\theta = p/5 = 0.01$ are identical for $N = 100$ because the increments in h are $2/N = 0.02$ and thus these two thresholds are indistinguishable for this value of N for discrete h .

priately selected firing threshold, for all other choices merely exacerbate the differences, not remove them. We may, however, put the question the other way around: what value of SNR, other than unity, would be necessary in order to match SNR memory lifetimes to MFPT memory lifetimes in the logarithmic asymptotic limit (necessarily with $\theta = 0$)? If we take a target SNR of ϑ rather than unity, then Eq. (6.18) matches Eq. (6.14) when $\vartheta^2 = \exp(-\hat{\gamma})/2$, or when $\vartheta \approx 0.530$. That is, the mean memory signal must reach zero at nearly the half a standard deviation level if $\tau_{\text{snr}}(p)$ is to agree with $\tau_{\text{mfpt}}(p)$. Needless to say, such a loose criterion seems somewhat unreasonable: usually, we would insist on increasing the number of standard deviations, not decreasing them.

9 Discussion

Several approaches to determining the lifetimes of memories in feed-forward or recurrent networks of neurons with binary synapses have been considered, including the signal-to-noise ratio (Tsodyks, 1990; Amit and Fusi, 1994; Brader et al., 2007), signal detection theory (Leibold and Kempster, 2006, 2008) and retrieval probabilities (Huang and Amit, 2010, 2011). To our knowledge, our analysis above is the first time that an approach based on mean first passage times for a perceptron’s output to fall below firing threshold has been considered. The difference between SNRs and MFPTs in this memory lifetime context may be compared to the different ways in which to compute the time-to-spike in a leaky integrate-and-fire neuron (see, for example, Tuckwell, 1988;

Gerstner and Kistler, 2002). Although considerably harder than an SNR calculation, a time-to-spike calculation based on the first passage time for the membrane potential to reach threshold takes into account all possible fluctuations that drive the membrane potential to threshold. Furthermore, there is no arbitrariness in an MFPT calculation as there is in an SNR calculation in defining when the mean signal reaches threshold at a certain, fixed number of standard deviations. Moreover, unlike SNRs, MFPTs are identical in discrete-time and continuous-time formulations.

We have studied extensively above the MFPT memory lifetime for the simplest possible, memory-related synaptic plasticity model based on a synapse without internal structure (Tsodyks, 1990). The use of such a simple model has greatly facilitated the derivation of exact or approximate results that would, in a more complicated model, be intractable. The continuum h limit naturally presents itself for N large enough, and for the update probability p sufficiently small, an OU limit essentially falls out. Typically a diffusion approximation is necessary in time-to-spike calculations, also resulting in an OU process for the membrane potential dynamics. It should not be surprising, therefore, that a limiting OU process appears in our analysis of MFPT memory lifetimes. On such a view, the initial memory encoding event pulls out a spring by an amount determined by the strength of the initial encoding. Ongoing memory storage constitutes a restoring force that pulls the spring back to equilibrium, with oscillations that asymptote to a fixed amplitude in the limit of large time.

Although the OU limit is an approximation, one that ignores the potentially

large jump processes that are present in the dynamics of the memory signal h , it provides qualitatively very good understanding of the full numerical results based on the Gaussian kernel or the controlled approximations based on the formal or Laplacian kernels. The results for the formal kernel asymptote to those for the OU process because the formal kernel has zero width and therefore explicitly lacks jump processes. In some sense, then, the formal kernel is a generalisation of the OU process but still based on a diffusion approximation: the differential equation governing $\tau_{\text{mfpt}}(x)$ for the formal kernel in Eq. (7.5) is second order. Conversely, the Laplacian kernel retains jump processes and provides extremely good quantitative agreement with the numerical results for the Gaussian kernel. Although Eq. (7.15) is superficially also second order, in fact the $\tau_{\text{mfpt}}(qx)$ term on the right hand side in effect induces derivatives in x of all orders for $p \neq 0$. It cannot, therefore, arise from a purely second-order Fokker-Planck equation. To offset this good agreement, however, the analytical results for the Laplacian kernel are somewhat complicated, precisely because this $\tau_{\text{mfpt}}(qx)$ term induces an expansion at all orders in powers of p .

Despite the fact that the full solutions for $\tau_{\text{mfpt}}(x)$ for the Laplace kernel and even for the formal kernel are complicated, fortunately their asymptotic behaviours in the large Nx^2 regime are considerably simpler, with all results at leading order in $1/p$ necessarily matching the asymptotic behaviour for the OU limit. This behaviour (at least for $\theta = 0$) is logarithmic in Nx^2 , so that MFPT memory lifetimes grow only logarithmically as a function of the number of synapses, N , when N is large enough. An identical result is of course by

now classic for SNR memory lifetimes in palimpsests, and much work has been devoted to overcoming the problems caused by such absurdly inefficient memory storage. However, there are significant and important differences between the MFPT and SNR results.

First, the MFPT result is only an asymptotic, large Nx^2 result, while the SNR result is typically regarded as generic. It is true that the SNR result is coupled with the requirement that $Np^2 > 1$ in order to achieve an initial SNR in excess of unity. This is because of internal consistency, in the sense that SNR memory lifetimes are defined as the time at which the SNR drops to unity, and so the initial SNR must be in excess of unity. Thus, the requirement that $Np^2 > 1$ simply arises rather trivially from the very definition of SNR lifetimes. It emphatically does not follow from a mathematical, large Nx^2 argument with $x = p$ in order to realise the asymptotic behaviour of some function. It is also worth pointing out that the motivation, other than definitional consistency, for demanding an initial SNR in excess of unity is somewhat unclear. Some authors have argued that a large initial SNR correlates with a well-encoded, “vivid” memory. However, the only relevant read-out from a neuron is its thresholded output, here of $+1$ or -1 : it does not matter how small its total input is provided this total input exceeds firing threshold, because the Heaviside non-linearity takes care of the rest. The vividness of the memory will correlate with the neuron’s firing, not with its sub-threshold membrane integration dynamics. If we simply insist that the initial memory is well-encoded, then we require only that $Np \gg 1$, i.e. that many synapses changed due to its storage. The two

conditions, $Np \gg 1$ and $Np^2 > 1$, may be rather different: consider, for example, the threshold case $p = 1/\sqrt{N}$ with N large enough.

Second, extending the asymptotic results outside of their range of validity, we appear to obtain lower bounds on the minimum number of synapses required to achieve positive memory lifetimes. Specifically, for MFPTs in the OU limit, we saw earlier that we require, approximately, $N > 1/4p^2$. With SNRs, we have the even stronger form, $N > 1/p^2$, i.e. four-fold more synapses are apparently required. However, MFPT memory lifetimes are positive for all non-zero N : there is no lower bound on N . It is true that for very small N , the initial memory encoding may be weak, especially if p is small, but such considerations are already factored into the MFPT calculation. Specifically, for small N , some neurons' total inputs upon initial memory storage will not exceed firing threshold and thus their first passage times will be identically zero; other neurons' total inputs will exceed firing threshold and contribute to a small but distinctly non-zero first passage time. Despite these marginal dynamics for small N , we saw in Fig. 4 that where the asymptotic MFPT or SNR results massively over-estimate the minimum number of synapses for non-zero memory lifetimes, the actual memory lifetimes for smaller numbers of synapses than this apparent minimum are in fact a sizeable fraction of the lifetimes for very much larger values of N for which the asymptotic behaviour is valid. That this is possible is due to the very slow, logarithmic growth: prior to the onset of this logarithmic growth, the memory lifetimes will already have grown to a considerable fraction of the lifetimes at logarithmic onset. For

example, in Fig. 4D at $N = 2000$, below the supposed lower bound on N from the asymptotic behaviour and well below that from SNR, we already have $r\tau_{\text{mfpt}}(p) \approx 60$, while at $N = 10^5$, 50-fold larger than $N = 2000$, we still only have $r\tau_{\text{mfpt}}(p) \approx 195$, barely 3-fold higher than at $N = 2000$. Because of this massive over-estimate of the minimum number of synapses required for memory storage, basing arguments about network function on this asymptotic, logarithmic behaviour may be profoundly dangerous, leading to false conclusions about any particular model's viability, or absence thereof, as a model of memory. This conclusion is true also of the logarithmic, SNR behaviour.

In the OU limit, the extension of our analysis to include a non-zero firing threshold θ was immediate, since it merely involved a translation of the absorbing boundary. For $\theta = 0$, we have the usual asymptotic, logarithmic behaviour of $\tau_{\text{mfpt}}(x)$ for large Nx^2 . However, for any value of $\theta > 0$ we found the remarkable property that $\tau_{\text{mfpt}}(x) \sim \frac{1}{p} \log_e \left(\frac{x}{\theta} \right)$ for large enough N , i.e. $\tau_{\text{mfpt}}(x)$ becomes *entirely independent of N* . Thus, MFPT memory lifetimes in the presence of a positive firing threshold are strictly bounded from above as a function of N and do not grow logarithmically with N . We confirmed this analytical observation, based on the OU limit, by comparing it to MFPTs from simulations away from the OU limit, seeing precisely the same asymptoting of MFPTs as a function of N for positive firing thresholds. While the feeble logarithmic growth of memory lifetimes from SNRs and asymptotic MFPTs is typically regarded as seriously problematic for models of memory that aspire to be based on biologically plausible principles, this strict upper bound

on memory lifetimes for positive firing thresholds must surely be regarded as a catastrophe. Just as the introduction of internal, plasticity-related synaptic states has been proposed in an attempt to overcome or at least ameliorate the logarithmic growth of memory lifetimes (see, for example, Fusi et al., 2005; Leibold and Kempner, 2008; Elliott and Lagogiannis, 2012), we might hope that such models may also overcome or soften this strict bounding of memory lifetimes in the presence of a positive firing threshold. However, at least in the case of cascade-based synapses (Fusi et al., 2005) and filter-based synapses (Elliott and Lagogiannis, 2012), preliminary simulation-based evidence does not indicate this to be the case: MFPT memory lifetimes still appear to asymptote to a constant, fixed value independent of N for a positive firing threshold.

It is natural, then, to wonder what is so special about a firing threshold of $\theta = 0$, which permits logarithmic growth in memory lifetimes, as opposed to $\theta > 0$, which strictly bounds memory lifetimes from above. From a purely technical standpoint, the result that $r\tau_{\text{mfpt}}(x) \sim \frac{1}{p} \log_e \left(\frac{x}{\theta} \right)$ in Eq. (6.16) requires that it is possible to make $N\theta^2$ (strictly speaking, $\sqrt{N}\theta$) large. Clearly, for any positive value of θ , regardless of how close to zero, N may always be chosen so that $N\theta^2$ is large. Conversely, for $\theta = 0$, it is never possible to make $N\theta^2$ large because its value is always identically zero. By taking θ small enough but non-zero, we can make $\tau_{\text{mfpt}}(x)$ grow (very nearly) logarithmically over some (in principle large) range of N , but increasing N beyond that range will always cause $\tau_{\text{mfpt}}(x)$ to peel off from logarithmic growth and asymptote to its upper bound. Of course, the onset of this asymptotic behaviour may occur for values

of N that are too large to be biologically relevant. Nevertheless, formally, the $\theta \rightarrow 0$ limit is discontinuous and results at $\theta = 0$ are untypical and non-generic: the behaviours for $\theta > 0$ and $\theta = 0$ are qualitatively different.

But why $\theta = 0$ specifically and not some other value of the neuron’s firing threshold? The reason for this is simply that for the way in which we have set up the memory encoding problem in the above, $\mathbb{E}[h(t)] \rightarrow 0$ as $t \rightarrow \infty$, i.e. the mean memory signal asymptotes to zero and not some other value. The mean memory signal asymptotes to zero because we have employed binary synapses with strengths ± 1 ; symmetric, zero-mean inputs through these synapses; and symmetric potentiation and depression processes. These are the standard, simplifying assumptions classically made when gauging memory lifetimes in models such as those considered here. However, were we to relax any of these assumptions, $\mathbb{E}[h(t)]$ would typically asymptote to some non-zero value, call it h_∞ . We could then work with $h(t) - h_\infty$ above, in effect performing a transformation that would restore the standard assumptions. For example, if we added an overall constant, say σ , to the ± 1 synaptic strengths considered here, then our entire analysis above would simply translate h by this overall constant and we would have $h_\infty = \sigma$. Thus, in general, the “critical” firing threshold would be $\theta = h_\infty$ and not $\theta = 0$. For $\theta = h_\infty$, the logarithmic growth of memory lifetimes would be observed, while for $\theta > h_\infty$, the logarithmic behaviour would be destroyed.

In comparing memory lifetimes from SNRs to those from MFPTs, it may be argued that we are not comparing like-with-like, and thus qualitative differ-

ences are to be expected. Is it possible, for example, that by making the SNR threshold or the neuron's firing threshold scale with the system size, here determined by the number of synapses N , we could restore equivalence between the two measures? If we were to scale the required SNR threshold as some power of N , say ϑN^β , then memory lifetimes would grow as $(\frac{1}{2} - \beta) \log_e N$, and for the particular choice $\beta = \frac{1}{2}$, the logarithmic growth would drop out, achieving equivalence with the MFPT form for $\theta \neq 0$. Notice that again this would appear to be a generic result, although in fact we know that it must only be an asymptotically valid, large N result. Conversely, if we were to scale the neuron's firing threshold θ with the system size, writing instead $\theta = \bar{\theta}/N^\beta$ with $\beta > 0$, so that the firing threshold drops as N increases, then immediately the asymptotic result $r\tau_{\text{mfpt}}(x) \sim \frac{1}{p} \log_e \left(\frac{x}{\theta}\right)$ would regain a logarithmic dependence on N with overall coefficient β , provided that $\beta < \frac{1}{2}$. This proviso is necessary because Eq. (6.16) is valid only if $N\theta^2$ can be made large. If $\theta = \bar{\theta}/N^\beta$, then $N\theta^2 = N^{1-2\beta} \bar{\theta}^2$, and for $\beta \geq \frac{1}{2}$, making $N\theta^2$ large by increasing N is not possible. However, in this case, $\tau_{\text{mfpt}}(\theta)$ simply never attains its asymptotic, logarithmic behaviour, and so never cancels that from $\tau_{\text{mfpt}}(x)$ in Eq. (6.13). Thus, regardless of the value of $\beta > 0$, logarithmic behaviour in the MFPT is re-instated by taking $\theta = \bar{\theta}/N^\beta$. We would have to set $\beta = \frac{1}{2}$ to achieve equivalence with the SNR form.

Neither move, however, appears justified. First, the entire basis of the enhancement of memory lifetimes, gauged by SNRs, is that as the number of synapses increases, the law of large numbers dominates: the fluctuations are

suppressed since they go like $1/\sqrt{N}$ (ignoring any covariance-induced terms) while the signal remains unchanged, not scaling with N . Thus, we would expect memory lifetimes to increase as N increases, ultimately because $h(t)$ is a sum over N identically-distributed random variables and the signal stands out more from the background noise as N increases. While logarithmic growth is theoretically optimal for the simple memory model considered here, slower growth, asymptoting to a constant as $N \rightarrow \infty$, is more usual. Second, scaling the neuron’s firing threshold with N is invalid for two reasons. The first reason is that the neuron’s output $h(t)$ does not scale with N (or, rather, in including a factor of $1/N$ in the definition of $h(t)$, we have already in effect scaled the threshold correctly), and increasing N will therefore not change $\mathbb{E}[h(t)]$ and so there is no need to scale θ as N increases. The second reason is that scaling θ to re-introduce a logarithmic dependence on N only works “naturally” for the specific, critical value of $h_\infty = 0$. For $h_\infty \neq 0$, we would in fact have to scale only that part of the neuron’s firing threshold that is above h_∞ , i.e. we would have to write $\theta = h_\infty + \bar{\theta}/N^\beta$. However, this is unnatural, not least because a neuron cannot “know” in advance what h_∞ is, because it does not have access to the inputs’ statistics, nor arguably to the possible distribution of synaptic strengths, given contaminating factors such as noise, etc.

Why, then, is there this difference between SNR and MFPT memory lifetimes for $\theta = 0$ (or $\theta = h_\infty$) and $\theta > 0$ (or $\theta > h_\infty$)? This difference arises not because of a failure to scale SNR or firing thresholds correctly with system size, but because the MFPT relates directly to *observable*, supra-firing-

threshold memory performance while the SNR constitutes a “bird’s eye view” that may, and typically almost certainly will, involve dynamics that are sub-firing-threshold. It is for precisely this reason that the MFPT measure is to be preferred over the SNR measure. The SNR measure gives theoretically optimal memory performance without indicating whether memories may actually be read out. The MFPT measure gives an in-practice, realisable memory performance because the stored memories can, by definition, always be read out. The key issue determining the in-practice realisability of optimal, logarithmic memory performance is whether or not we have the freedom to set the neuron’s firing threshold appropriately. If we always have the freedom to set $\theta = h_\infty$, then we can always realise optimal memory performance with logarithmic growth, and in this case, as we might expect and as is indeed the case, SNRs and MFPTs agree, at least asymptotically for large N , although we have seen that SNRs away from this asymptotic regime are unreliable. But consider the extreme scenario in which $\theta = \mathbb{E}[h_0]$, so that the firing threshold sits precisely at the (average) strength of the initial encoding. Clearly, the sub-firing-threshold $h(t)$ dynamics are entirely unaffected by the location of the firing threshold, and thus $h(t)$ will still fall exponentially fast, asymptotically to h_∞ , and with fluctuations in $h(t)$ being suppressed as N increases. The neuron’s sub-firing-threshold dynamics (the neuron’s integrative membrane potential dynamics) retain a trace of the tracked memory. But, critically, this trace is inaccessible, in terms of normal memory performance, because it is entirely sub-firing-threshold. The trace may be made accessible only by re-

ducing the neuron's firing threshold. The extent to which we reduce the firing threshold determines the relative attainability of that threshold to the ongoing fluctuations that will degrade the memory. If it is reduced all the way to h_∞ , then we achieve optimal, logarithmic memory performance; but if to some value between h_∞ and $\mathbb{E}[h_0]$, then the memory performance will be sub-optimal and so sub-logarithmic and in fact asymptotically bounded. The firing threshold $\theta = h_\infty$ is special, and privileged, because $\mathbb{E}[h(t)]$ for this specific choice always remains above the firing threshold. However, for $\theta > h_\infty$, $\mathbb{E}[h(t)]$ will always eventually drop below firing threshold. It is this dropping below firing threshold that causes the logarithmic or near-logarithmic enhancement of memory lifetimes with increasing N to be truncated, because the memory signal, now hiding below the firing threshold, becomes inaccessible.

We may see this truncation explicitly by considering a modified form of an SNR calculation. Rather than requiring that $\mathbb{E}[h(\tau_{\text{snr}}(x))] = 1/\sqrt{N}$ to determine $\tau_{\text{snr}}(x)$ (with $h_\infty = 0$), we instead enforce an accessibility criterion by requiring that the mean memory signal reaches one standard deviation above the neuron's firing threshold. Thus, we need $\mathbb{E}[h(\tau_{\text{snr}'}(x))] = \theta + 1/\sqrt{N}$, whence $pr\tau_{\text{snr}'}(x) = \log_e x/(\theta + 1/\sqrt{N})$. Only for $\theta > 0$ (or more generally $\theta > h_\infty$) may we now perform a Taylor expansion around $1/N = 0$, obtaining $r\tau_{\text{snr}'}(x) \sim \frac{1}{p} \log_e \left(\frac{x}{\theta}\right) - \frac{1}{p\theta\sqrt{N}}$. Thus, the logarithmic growth is indeed truncated in an SNR-type calculation for $\theta > 0$ (but not for $\theta = 0$ because the Taylor expansion ceases to be valid), and we see the asymptotic bounding of memory lifetimes in agreement with the MFPT calculation. This agreement is achieved

precisely by imposing an accessibility or memory read-out criterion when $\theta > 0$, essentially throwing away any inaccessible and therefore irrelevant sub-firing-threshold dynamics.

In an artificial network or in electronic hardware, we always have the freedom to pick a neuron's firing threshold and thus to realise theoretically-optimal, logarithmic memory performance. However, if models such as those considered above have any relation, for example, to memory performance in real neurons (and, after all, the motivation for considering bounded synaptic strengths is increased biological plausibility), then the freedom to set firing thresholds at will to achieve optimal memory performance is likely not available. In real neurons, then, we might expect logarithmic memory performance to be untypical and non-generic, even if in-principle achievable with suitably-set firing thresholds. These considerations clearly apply not only to synapses without internal structure but also to those with internal structure, which is presumably why our preliminary simulation-based evidence mentioned above indicates that we also observe a breaking of the logarithmic growth in these more complicated models.

Analytical extension of our MFPT approach to memory lifetimes for these more complicated synapse models is likely to be hard. The non-Markovian nature of the transitions in h in the presence of synapses with internal structure means that we must work in principle with the full product space, flattened or otherwise, working at the level of individual synaptic strength transitions and not transitions in h . Numerically, working with such large matrices will

rapidly become intractable for all but the simplest problems with very small N . One promising approach, however, is to attempt to develop a stochastic updater synapse approximation to these more complicated synapse models. For the transitions in h , we are not interested in the transitions in the internal states of the synapses. Instead, we need only know when the synapses change strength. Thus, we could integrate out the internal, plasticity-related states as we have done in other analyses (Elliott, 2010, 2011b), and consider a structureless, stochastically-updating synapse in which the probability p for changing strength in response to an induction stimulus becomes both time- and strength-state-dependent, yielding $p^\pm(t)$. Such an approach would reproduce exactly the dynamics of the mean memory signal, essentially by construction. The h transition matrix \mathbb{T} will become time- or time-step-dependent and thus simple powers of \mathbb{T} will no longer suffice for multi-step transitions. Instead we will have to deal with products of \mathbb{T} evaluated at different time steps. Nevertheless, we would need work only with $(N + 1) \times (N + 1)$ -dimensional matrices. Integrating out internal states will inevitably mean that we do not fully capture all the fluctuations at all orders that lead to changes in synaptic strength, changes in h , and therefore first passage events past the firing threshold θ . However, the approximation may be good enough to provide qualitative understanding of MFPT memory lifetimes with more complicated synapses, and it will be interesting to pursue this approach in subsequent work.

Acknowledgements: I thank one of the reviewers for suggesting the system-

size argument, although I believe that the core of the difference between MFPT and SNR memory performance is in-practice versus in-principle realisability, as discussed above.

Appendix A Non-Poisson Continuous-Time Processes

The equality of the discrete-time and continuous-time MFPTs, up to an overall rate factor, arises because in the derivation of Eq. (3.9) for the continuous-time MFPT, we employed the result

$$\int_0^\infty dt \frac{(rt)^m}{m!} e^{-rt} = \frac{1}{r}, \quad (\text{A.1})$$

where the integrand is just the Poisson probability for the occurrence of precisely m events in time t . In fact, this result, suitably generalised, carries over to any continuous-time process in which the waiting times between memory storage events are governed by any common probability density function and not just the exponentially-distributed waiting times that define the Poisson process. We show this here.

Let the probability density function governing the waiting times be $f(t)$ and let $\pi_m(t)$ be the probability of precisely m events in time t . Let $\bar{F}(t)$ be the probability of the non-occurrence of a single event in time t . Then $\pi_m(t)$

is given by the convolution

$$\begin{aligned} \pi_m(t) = & \int_0^t dt_1 f(t_1) \int_0^{t-t_1} dt_2 f(t_2) \cdots \int_0^{t-t_1-\cdots-t_{m-1}} dt_m f(t_m) \\ & \times \int_0^{t-t_1-\cdots-t_m} dt_{m+1} \bar{F}(t_{m+1}), \end{aligned} \quad (\text{A.2})$$

where the last integral enforces the non-occurrence of an event in the time between the occurrence of the m -th event and time t . If a function with a hat over it denotes its Laplace transform and s is the transformed variable, then

$$\hat{\pi}_m(s) = [\hat{f}(s)]^m \hat{F}(s), \quad (\text{A.3})$$

where $\hat{F}(s) = [1 - \hat{f}(s)]/s$. The equivalent of the integral in Eq. (A.1) is just $\int_0^\infty dt \pi_m(t) \equiv \hat{\pi}_m(0)$, where if necessary we define $\hat{\pi}_m(0) = \lim_{s \rightarrow 0} \hat{\pi}_m(s)$. Since $f(t)$ is a probability density function, we have that $\int_0^\infty dt f(t) \equiv \hat{f}(0) = 1$. So,

$$\int_0^\infty dt \pi_m(t) = \lim_{s \rightarrow 0} \left\{ [\hat{f}(s)]^m \frac{1 - \hat{f}(s)}{s} \right\} = -\hat{f}'(0), \quad (\text{A.4})$$

where $\hat{f}'(s) = d\hat{f}(s)/ds$. In fact, $-\hat{f}'(0)$ is just the reciprocal of the asymptotic rate of the overall, continuous-time process. The rate $r(t)$ of this overall process is defined as

$$r(t) = \frac{d}{dt} \mathbb{E}[N(t)], \quad (\text{A.5})$$

where $\mathbb{E}[N(t)]$ is the expectation value of the number of events $N(t)$ in time t , so that $\mathbb{E}[N(t)] = \sum_{m=0}^\infty m \pi_m(t)$. It is easy to see that

$$\hat{r}(s) = \frac{\hat{f}(s)}{1 - \hat{f}(s)}. \quad (\text{A.6})$$

The asymptotic rate $\lim_{t \rightarrow \infty} r(t)$ is just $\lim_{s \rightarrow 0} s \hat{r}(s)$ (see, for example, Elliott, 2010), and thus

$$\lim_{t \rightarrow \infty} r(t) = \lim_{s \rightarrow 0} \frac{s \hat{f}(s)}{1 - \hat{f}(s)} = -\frac{1}{\hat{f}'(0)}. \quad (\text{A.7})$$

So,

$$\int_0^\infty dt \pi_m(t) = \frac{1}{\lim_{t \rightarrow \infty} r(t)}, \quad (\text{A.8})$$

and therefore continuous-time and discrete-time MFPTs are always identical, up to this overall asymptotic rate factor, for any common waiting time distribution $f(t)$.

Appendix B Computation of \mathbb{T}^m

By defining the generating function

$$\mathcal{G}_1(x|k) = \sum_{i=0}^N x^i \text{Prob} \left[h_{n+1} = \frac{2i}{N} - 1 \mid h_n = \frac{2k}{N} - 1 \right], \quad (\text{B.1})$$

for the elements of \mathbb{T} given in Eq. (5.2), we find that

$$\mathcal{G}_1(x|k) = \left[\frac{p}{2} + \left(1 - \frac{p}{2}\right) x \right]^k \left[\left(1 - \frac{p}{2}\right) + \frac{p}{2} x \right]^{N-k}. \quad (\text{B.2})$$

From this we may obtain the generating function for the elements of \mathbb{T}^m . For \mathbb{T}^2 , for example, we have that

$$\begin{aligned} \mathcal{G}_2(x|k) &= \sum_{i=0}^N x^i \sum_{j=0}^N \text{Prob} \left[h_{n+2} = \frac{2i}{N} - 1 \mid h_n = \frac{2j}{N} - 1 \right] \\ &\quad \times \text{Prob} \left[h_{n+1} = \frac{2j}{N} - 1 \mid h_n = \frac{2k}{N} - 1 \right] \\ &= \sum_{j=0}^N \mathcal{G}_1(x|j) \text{Prob} \left[h_{n+1} = \frac{2j}{N} - 1 \mid h_n = \frac{2k}{N} - 1 \right] \\ &= \left\{ \frac{1}{2} [1 - (1-p)^2] + \frac{1}{2} [1 + (1-p)^2] x \right\}^k \\ &\quad \times \left\{ \frac{1}{2} [1 + (1-p)^2] + \frac{1}{2} [1 - (1-p)^2] x \right\}^{N-k}. \quad (\text{B.3}) \end{aligned}$$

In general, by induction, we find that \mathbb{T}^m has elements given by the generating function

$$\mathcal{G}_m(x|k) = \left[\frac{1}{2}(1 - q^m) + \frac{1}{2}(1 + q^m)x \right]^k \left[\frac{1}{2}(1 + q^m) + \frac{1}{2}(1 - q^m)x \right]^{N-k}, \quad (\text{B.4})$$

where $q = 1 - p$. The unconditional generating function for the distribution of h_m , $\mathcal{G}_m(x)$, is then given by

$$\begin{aligned} \mathcal{G}_m(x) &= \sum_{k=0}^N \mathcal{G}_m(x|k) \text{Prob} \left[h_0 = \frac{2k}{N} - 1 \right] \\ &= \left[\frac{1}{2}(1 - pq^m) + \frac{1}{2}(1 + pq^m)x \right]^N. \end{aligned} \quad (\text{B.5})$$

We may obtain results in continuous time rather than discrete time by writing the continuous-time generating functions as

$$\mathcal{G}(t; x|k) = e^{-rt} \sum_{m=0}^{\infty} \frac{(rt)^m}{m!} \mathcal{G}_m(x|k), \quad (\text{B.6})$$

$$\mathcal{G}(t; x) = e^{-rt} \sum_{m=0}^{\infty} \frac{(rt)^m}{m!} \mathcal{G}_m(x). \quad (\text{B.7})$$

Unfortunately, we cannot evaluate these sums explicitly without expanding $\mathcal{G}_m(x|k)$ and $\mathcal{G}_m(x)$ using the binomial theorem, producing messy results. Nevertheless, we can obtain the moments of the continuous-time process in the standard way by differentiation with respect to x , and then the Poisson sums do simplify.

Appendix C Derivation of Eq. (7.22)

In order to derive Eq. (7.22), we need the q -binomial theorem and some of its corollaries. A discussion of q -series and the q -binomial theorem may be found,

for example, in Andrews et al. (1999). The q -binomial theorem states that, for $|x| < 1$ and $|q| < 1$,

$$\sum_{k=0}^{\infty} \frac{(a; q)_k}{(q; q)_k} x^k = \frac{(ax; q)_{\infty}}{(x; q)_{\infty}}. \quad (\text{C.1})$$

Four of its corollaries are

$$\sum_{k=0}^n \begin{bmatrix} n \\ k \end{bmatrix}_q (-1)^k q^{kC_2} x^k = (x; q)_n, \quad (\text{C.2})$$

$$\sum_{k=0}^{\infty} \begin{bmatrix} n+k-1 \\ k \end{bmatrix}_q x^k = \frac{1}{(x; q)_n}, \quad (\text{C.3})$$

$$\sum_{k=0}^{\infty} \frac{(-1)^k q^{kC_2} x^k}{(q; q)_k} = (x; q)_{\infty}, \quad (\text{C.4})$$

$$\sum_{k=0}^{\infty} \frac{x^k}{(q; q)_k} = \frac{1}{(x; q)_{\infty}}, \quad (\text{C.5})$$

where the q -binomial coefficient $\begin{bmatrix} n \\ k \end{bmatrix}_q$ is given by

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \frac{(q; q)_n}{(q; q)_k (q; q)_{n-k}} = \begin{bmatrix} n \\ n-k \end{bmatrix}_q. \quad (\text{C.6})$$

We may now proceed to show that

$$\sum_{j=0}^{\infty} x q^{2j} (x; q^2)_j = 1 - (x; q^2)_{\infty}, \quad (\text{C.7})$$

from which Eq. (7.22) follows by putting $x = q^i$. We have

$$\begin{aligned} \sum_{j=0}^{\infty} x q^{2j} (x; q^2)_j &= \sum_{j=0}^{\infty} x q^{2j} \sum_{k=0}^j (-x)^k (q^2)^{kC_2} \begin{bmatrix} j \\ k \end{bmatrix}_{q^2} && \text{by Eq. (C.2)} \\ &= - \sum_{k=0}^{\infty} (-x)^{k+1} (q^2)^{kC_2} \sum_{j=0}^{\infty} q^{2(j+k)} \begin{bmatrix} k+j \\ k \end{bmatrix}_{q^2} \\ &= - \sum_{k=0}^{\infty} (-x)^{k+1} (q^2)^{kC_2} q^{2k} \frac{1}{(q^2; q^2)_{k+1}} && \text{by Eq. (C.3)} \\ &= 1 - \sum_{l=0}^{\infty} (-x)^l (q^2)^{lC_2} \frac{1}{(q^2; q^2)_l} \\ &= 1 - (x; q^2)_{\infty}, && \text{by Eq. (C.4)} \end{aligned}$$

which establishes the advertised result.

Appendix D $\mathcal{O}(p^2)$ Expansion of $\tau_{\text{mfpt}}(x)$ for Laplace Kernel

We have the exact power series solution for $\tau_{\text{mfpt}}(x)$ for the Laplace $K_2(y|x)$ kernel in Eq. (7.23), given by

$$\tau_{\text{mfpt}}(x) = 1 + \frac{(q^2; q^2)_\infty}{(q; q^2)_\infty} [1 + g(x)] + f(x) \quad (\text{D.1})$$

where the particular solution $f(x)$ and the non-constant general solution $g(x)$ are

$$f(x) = - \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+2}}{(2m+2)!} (q^2; q^2)_m, \quad (\text{D.2})$$

$$g(x) = + \sum_{m=0}^{\infty} \frac{(q\gamma x)^{2m+1}}{(2m+1)!} (q; q^2)_m. \quad (\text{D.3})$$

We perform an expansion in powers of p . Expanding up to $\mathcal{O}(p^2)$, we write

$$f(x) = -\frac{N}{p} \left(1 - \frac{3}{2}p + \frac{1}{4}p^2 \right) x^2 \sum_{m=0}^{\infty} \frac{(Nx^2)^m}{(2m+2)!} c_m, \quad (\text{D.4})$$

where

$$c_m = 2^m m! \left[1 - p \frac{m(m+3)}{2} + p^2 \frac{m(9m^3 + 58m^2 + 69m - 64)}{72} \right]. \quad (\text{D.5})$$

Writing

$$f(x) = \frac{1}{p} [f_0(x) + p f_1(x) + p^2 f_2(x) + \dots], \quad (\text{D.6})$$

we find

$$f_0(x) = -\frac{Nx^2}{2} {}_2F_2\left(\{1, 1\}; \left\{\frac{3}{2}, 2\right\}, \frac{Nx^2}{2}\right), \quad (\text{D.7})$$

$$f_1(x) = \frac{Nx^2}{8} + \frac{\sqrt{2\pi N}}{16}x(3 + Nx^2)e^{Nx^2/2}\text{erf}\left(\sqrt{\frac{N}{2}}x\right) + \frac{Nx^2}{4} {}_2F_2\left(\{1, 1\}; \left\{\frac{3}{2}, 2\right\}, \frac{Nx^2}{2}\right) \quad (\text{D.8})$$

$$f_2(x) = -\frac{41}{96}Nx^2 + \frac{\sqrt{2\pi N}}{576}x(173 - 123Nx^2)e^{Nx^2/2}\text{erf}\left(\sqrt{\frac{N}{2}}x\right) - \frac{43}{144}Nx^2 {}_2F_2\left(\{1, 1\}; \left\{\frac{3}{2}, 2\right\}, \frac{Nx^2}{2}\right) - \frac{29}{432}(Nx^2)^2 {}_4F_4\left(\{2, 2, 2, 2\}; \left\{1, 1, \frac{5}{2}, 3\right\}, \frac{Nx^2}{2}\right) - \frac{1}{96}(Nx^2)^2 {}_5F_5\left(\{2, 2, 2, 2, 2\}; \left\{1, 1, 1, \frac{5}{2}, 3\right\}, \frac{Nx^2}{2}\right), \quad (\text{D.9})$$

where erf is the standard error function. Similarly for $g(x)$, we have

$$g(x) = \sqrt{\frac{N}{p}}\left(1 - \frac{3}{4}p - \frac{5}{32}p^2\right)x \sum_{m=0}^{\infty} \frac{(Nx^2)^m}{(2m+1)!}d_m, \quad (\text{D.10})$$

where

$$d_m = \frac{(2m-1)!}{2^{m-1}(m-1)!} \left[1 - p \frac{m(m+2)}{2} + p^2 \frac{m(9m^3 + 40m^2 + 18m - 49)}{72}\right], \quad (\text{D.11})$$

with $g_0 \equiv 1$. Then we write

$$g(x) = \frac{1}{\sqrt{p}} [g_0(x) + p g_1(x) + p^2 g_2(x) + \dots], \quad (\text{D.12})$$

and obtain

$$g_0(x) = \sqrt{\frac{\pi}{2}} \operatorname{erfi} \left(\sqrt{\frac{N}{2}} x \right), \quad (\text{D.13})$$

$$g_1(x) = -\frac{\sqrt{N} x (3 + Nx^2)}{8} e^{Nx^2/2} - \frac{3\sqrt{2\pi}}{16} \operatorname{erfi} \left(\sqrt{\frac{N}{2}} x \right), \quad (\text{D.14})$$

$$g_2(x) = \frac{5\sqrt{N}}{288} x \left[4(Nx^2)^2 + 13Nx^2 - 3 \right] e^{Nx^2/2} - \frac{5\sqrt{2\pi}}{96} \operatorname{erfi} \left(\sqrt{\frac{N}{2}} x \right) \\ + \frac{(Nx^2)^{3/2}}{48} {}_4F_4 \left(\left\{ \frac{3}{2}, 2, 2, 2 \right\}, \left\{ 1, 1, 1, \frac{5}{2} \right\}, \frac{Nx^2}{2} \right). \quad (\text{D.15})$$

From these expansions for $f(x)$ and $g(x)$, we may obtain that for $\tau_{\text{mfpt}}(x)$.

Appendix E Evaluation of Neumann Series for Laplace Kernel

For the Laplacian $K_2(y|x)$ kernel, we write the Neumann series is Eq. (7.3) as

$$\tau_{\text{mfpt}}(x) = \sum_{n=0}^{\infty} f_n(x), \quad (\text{E.1})$$

where

$$f_n(x) = \int_0^{\infty} dy f_{n-1}(y) K_2(y|x), \quad (\text{E.2})$$

for $n > 0$ and $f_0(x) \equiv 1$. By explicitly evaluating a few of these $f_n(x)$ for small n , we observe that we may write

$$f_n(x) = 1 + \sum_{i=1}^n \beta_i^{(n)} \exp(-\gamma q^i x), \quad (\text{E.3})$$

for some coefficients $\beta_i^{(n)}$. These coefficients may then be found to satisfy the set of recurrence relations,

$$\beta_i^{(n)} = \beta_1^{(n+1-i)} \frac{1}{(q^2, q^2)_{i-1}}, \quad (\text{E.4})$$

$$\beta_1^{(n+1)} = -\frac{1}{2} \left[1 + \sum_{i=1}^n \frac{1}{(1-q^i)} \frac{1}{(q^2; q^2)_{i-1}} \beta_1^{(n+1-i)} \right]. \quad (\text{E.5})$$

The quantities $\beta_1^{(n)}$ are fundamental because of Eq. (E.4). Writing a generating function for the $\beta_1^{(n)}$ as $g(z) = \sum_{n=0}^{\infty} \beta_1^{(n+1)} z^n$ and with the help of the auxiliary generating function

$$f(z) = \sum_{n=1}^{\infty} \frac{z^n}{(1-q^n)(q^2; q^2)_{n-1}}, \quad (\text{E.6})$$

we have, from Eq. (E.5),

$$g(z) = -\frac{1}{1-z} \frac{1}{2+f(z)}. \quad (\text{E.7})$$

We may in fact evaluate $f(z)$ explicitly using the methods in Appendix C:

$$\begin{aligned} \sum_{n=1}^{\infty} \frac{z^n}{(1-q^n)(q^2; q^2)_{n-1}} &= \sum_{l=0}^{\infty} (z q^l) \sum_{n=0}^{\infty} \frac{(z q^l)^n}{(q^2; q^2)_n} \\ &= \sum_{l=0}^{\infty} \frac{z q^l}{(z q^l; q^2)_{\infty}} \\ &= \frac{1}{(z; q^2)_{\infty}} \sum_{n=0}^{\infty} z q^{2n} (z; q^2)_n \\ &\quad + \frac{1}{(z q; q^2)_{\infty}} \sum_{n=0}^{\infty} (z q) q^{2n} (z q; q^2)_n \\ &= \frac{1 - (z; q^2)_{\infty}}{(z; q^2)_{\infty}} + \frac{1 - (z q; q^2)_{\infty}}{(z q; q^2)_{\infty}} \\ &= \frac{(z; q^2)_{\infty} + (z q; q^2)_{\infty}}{(z; q)_{\infty}} - 2. \end{aligned} \quad (\text{E.8})$$

Since $(z; q)_{\infty} / (1-z) = (z q; q)_{\infty}$, we have

$$g(z) = -\frac{(z q; q)_{\infty}}{(z; q^2)_{\infty} + (z q; q^2)_{\infty}}. \quad (\text{E.9})$$

We now write

$$\tau(x; z) = \sum_{n=0}^{\infty} z^n f_n(x) \quad (\text{E.10})$$

and we of course have

$$\lim_{z \rightarrow 1} \tau(x; z) = \tau_{\text{mfpt}}(x). \quad (\text{E.11})$$

By writing

$$\tau(x; z) = \begin{pmatrix} 1 & 1 & 1 & 1 & \dots \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 & \dots \\ z & z\beta_1^{(1)} & 0 & 0 & \dots \\ z^2 & z^2\beta_1^{(2)} & z^2\beta_2^{(2)} & 0 & \dots \\ z^3 & z^3\beta_1^{(3)} & z^3\beta_2^{(3)} & z^3\beta_3^{(3)} & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} 1 \\ e^{-\gamma q^1 x} \\ e^{-\gamma q^2 x} \\ e^{-\gamma q^3 x} \\ \vdots \end{pmatrix}, \quad (\text{E.12})$$

and using Eq. (E.4) we see that we have

$$\begin{aligned} \tau(x; z) &= \frac{1}{1-z} + z g(z) e^{-\gamma q^1 x} + \frac{z^2 g(z)}{(q^2; q^2)_1} e^{-\gamma q^2 x} + \frac{z^3 g(z)}{(q^2; q^2)_2} e^{-\gamma q^3 x} + \dots \\ &= \frac{1}{1-z} + z g(z) \sum_{n=0}^{\infty} \frac{z^n \exp(-\gamma q^{n+1} x)}{(q^2; q^2)_n}. \end{aligned} \quad (\text{E.13})$$

Although there appears to be a divergence as $z \rightarrow 1$ because of the $(1-z)^{-1}$ term on the right hand side, this is not in fact the case. Nevertheless, we remove this $(1-z)^{-1}$ term by considering only $\tau(x; z) - \tau(0, z)$. We have that

$$\tau(x; z) - \tau(0, z) = z g(z) \sum_{n=0}^{\infty} \frac{\exp(-\gamma q^{n+1} x) - 1}{(q^2; q^2)_n} z^n. \quad (\text{E.14})$$

Since

$$g(1) = -\frac{(q; q)_{\infty}}{(1; q^2)_{\infty} + (q; q^2)_{\infty}} = -\frac{(q; q)_{\infty}}{(q; q^2)_{\infty}} \equiv -(q^2; q^2)_{\infty}, \quad (\text{E.15})$$

we have

$$\tau_{\text{mfpt}}(x) - \tau_{\text{mfpt}}(0) = (q^2; q^2)_{\infty} \sum_{n=0}^{\infty} \frac{1 - \exp(-\gamma q^{n+1} x)}{(q^2; q^2)_n}. \quad (\text{E.16})$$

It only remains to show that $\lim_{z \rightarrow 1} \tau(0; z)$ gives the correct, known value of

$\tau_{\text{mfpt}}(0)$:

$$\begin{aligned}
\tau(0; z) &= \frac{1}{1-z} + z g(z) \sum_{n=0}^{\infty} \frac{z^n}{(q^2; q^2)_n} \\
&= \frac{1}{1-z} + z g(z) \frac{1}{(z; q^2)_{\infty}} \\
&= \frac{1}{1-z} \left[1 - z \frac{(z q; q^2)_{\infty}}{(z; q^2)_{\infty} + (z q; q^2)_{\infty}} \right] \\
&= \frac{1}{1-z} \left[\frac{(1-z)(z q; q^2)_{\infty} + (z; q^2)_{\infty}}{(z; q^2)_{\infty} + (z q; q^2)_{\infty}} \right] \\
&= \frac{(z q; q^2)_{\infty} + (z q^2; q^2)_{\infty}}{(z; q^2)_{\infty} + (z q; q^2)_{\infty}} \\
&\xrightarrow{z \rightarrow 1} \frac{(q; q^2)_{\infty} + (q^2; q^2)_{\infty}}{(q; q^2)_{\infty}} \\
&= 1 + \frac{(q^2; q^2)_{\infty}}{(q; q^2)_{\infty}} \\
&\equiv \tau_{\text{mfpt}}(0),
\end{aligned}$$

as required.

References

- Abraham, W.C., Robins, A. 2005. Memory retention – the synaptic stability versus plasticity dilemma. *Trends in Neurosciences*, **28**, 73–78.
- Amit, D.J., Fusi, S. 1994. Learning in neural networks with material synapses. *Neural Computation*, **6**, 957–982.
- Andrews, G.E., Askey, R., Roy, R. 1999. *Special Functions*. Cambridge: Cambridge University Press.
- Bagal, A.A., Kao, J.P.Y., Tang, C.-M., Thompson, S.M. 2005. Long-term potentiation of exogenous glutamate responses at single dendritic spines. *Proceedings of the National Academy of Sciences U.S.A.*, **102**, 14434–14439.
- Brader, J.M., Senn, W., Fusi, S. 2007. Learning real-world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Computation*, **19**, 2881–2912.
- Elliott, T. 2010. A non-Markovian random walk underlies a stochastic model of spike-timing-dependent plasticity. *Neural Computation*, **22**, 1180–1230.
- Elliott, T. 2011a. The mean time to express synaptic plasticity in stochastic, integrate-and-express models of synaptic plasticity induction. *Neural Computation*, **23**, 124–159.
- Elliott, T. 2011b. Stability against fluctuations: Scaling, bifurcations and

- spontaneous symmetry breaking in stochastic models of synaptic plasticity. *Neural Computation*, **23**, 674–734.
- Elliott, T., Lagogiannis, K. 2012. The rise and fall of memory in a model of synaptic integration. *Neural Computation*, **24**, 2604–2654.
- Fusi, S., Drew, P.J., Abbott, L.F. 2005. Cascade models of synaptically stored memories. *Neuron*, **45**, 599–611.
- Gerstner, W., Kistler, W. 2002. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge: Cambridge University Press.
- Hertz, J., Krogh, A., Palmer, R.G. 1991. *Introduction to the Theory of Neural Computation*. Redwood City, CA: Addison Wesley.
- Hopfield, J.J. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences U.S.A.*, **79**, 2554–2558.
- Huang, Y., Amit, Y. 2010. Precise capacity analysis in binary networks with multiple coding level inputs. *Neural Computation*, **22**, 660–688.
- Huang, Y., Amit, Y. 2011. Capacity analysis in multi-state synaptic models: A retrieval probability perspective. *Journal of Computational Neuroscience*, **30**, 699–720.
- Leibold, C., Kempster, R. 2006. Memory capacity for sequences in a recurrent network with biological constraints. *Neural Computation*, **18**, 904–941.

- Leibold, C., Kempster, R. 2008. Sparseness constrains the prolongation of memory lifetime via synaptic metaplasticity. *Cerebral Cortex*, **18**, 67–77.
- Montgomery, J.M., Madison, D.V. 2002. State-dependent heterogeneity in synaptic depression between pyramidal cell pairs. *Neuron*, **33**, 765–777.
- Montgomery, J.M., Madison, D.V. 2004. Discrete synaptic states define a major mechanism of synapse plasticity. *Trends in Neurosciences*, **27**, 744–750.
- Nadal, J.P., Toulouse, G., Changeux, J.P., Dehaene, S. 1986. Networks of formal neurons and memory palimpsests. *Europhysics Letters*, **1**, 535–542.
- O’Connor, D.H., Wittenberg, G.M., Wang, S.S.-H. 2005a. Dissection of bidirectional synaptic plasticity into saturable unidirectional process. *Journal of Neurophysiology*, **94**, 1565–1573.
- O’Connor, D.H., Wittenberg, G.M., Wang, S.S.-H. 2005b. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proceedings of the National Academy of Sciences U.S.A.*, **102**, 9679–9684.
- O’Connor, D.H., Wittenberg, G.M., Wang, S.S.-H. 2007. Timing and contributions of pre-synaptic and post-synaptic parameter changes during unitary plasticity events at CA3-CA1 synapses. *Synapse*, **61**, 664–678.
- Parisi, G. 1986. A memory which forgets. *Journal of Physics A: Mathematical and General*, **19**, L617–L620.

- Petersen, C.C.H., Malenka, R.C., Nicoll, R.A., Hopfield, J.J. 1998. All-or-none potentiation at CA3-CA1 synapses. *Proceedings of the National Academy of Sciences U.S.A.*, **95**, 4732–4737.
- Sobczyk, A., Svoboda, K. 2007. Activity-dependent plasticity of the NMDA-receptor fractional Ca^{2+} current. *Neuron*, **53**, 17–24.
- Tsodyks, M.V. 1990. Associative memory in neural networks with binary synapses. *Modern Physics Letters B*, **4**, 713–716.
- Tuckwell, H.C. 1988. *Introduction to Theoretical Neurobiology*. Vol. 2. Cambridge: Cambridge University Press.
- van Kampen, N.G. 1992. *Stochastic Processes in Physics and Chemistry*. Amsterdam: Elsevier.
- Watson, G.N. 1936. The final problem: An account of the mock theta functions. *Journal of the London Mathematical Society*, **11**, 55–80.
- Yasuda, R., Sabatini, B.L., Svoboda, K. 2003. Plasticity of calcium channels in dendritic spines. *Nature Neuroscience*, **6**, 948–955.