Dynamic Integrative Synaptic Plasticity Explains the Spacing Effect in the Transition from Short- to Long-Term Memory

Terry Elliott[1]

Department of Electronics and Computer Science,

University of Southampton, Highfield,

Southampton, SO17 1BJ,

United Kingdom.

June 19, 2019.

Running Title: Spacing effect.

[1]Tel: +44 (0)23 8059 6000, Fax: +44 (0)23 8059 2783, Email: te@ecs.soton.ac.uk.

Abstract

Repeated stimuli that are spaced apart in time promote the transition from short- to long-term memory, while massing repetitions together does not. Previously, we showed that a model of integrative synaptic plasticity, in which plasticity induction signals are integrated by a low-pass filter before plasticity is expressed, gives rise to a natural timescale at which to repeat stimuli, hinting at a partial account of this spacing effect. The account was only partial because the important role of neuromodulation was not considered. We now show that by extending the model to allow dynamic integrative synaptic plasticity, the model permits synapses to robustly discriminate between spaced and massed repetition protocols, suppressing the response to massed stimuli while maintaining that to spaced stimuli. This is achieved by dynamically coupling the filter decay rate to neuromodulatory signalling in a very simple model of the signalling cascades downstream from cAMP production. In particular, the model's parameters may be interpreted as corresponding to the duration and amplitude of the waves of activity in the MAPK pathway. We identify choices of parameters and repetition times for stimuli in this model that optimise the ability of synapses to discriminate between spaced and massed repetition protocols. The model is very robust to reasonable changes around these optimal parameters and times, but for large changes in parameters, the model predicts that massed and spaced stimuli cannot be distinguished, or that the responses to both patterns are suppressed. A model of dynamic integrative synaptic plasticity therefore explains the spacing effect under normal conditions, and also predicts its breakdown under abnormal conditions.

1 Introduction

Memory is a complex, multi-faceted phenomenon (Eichenbaum & Cohen, 2001) that has been studied in animals as diverse as *Aplysia* (Hawkins *et al.*, 2006), *Drosophila melanogaster* (Heisenberg, 2003; Davis, 2005) and mice, rats and humans (Andersen *et al.*, 2007). Synaptic plasticity is believed to underlie memory formation. The standard mammalian model of synaptic plasticity is long-term potentiation (LTP) (Bliss & Lømo, 1973) and long-term depression (LTD) (Lynch *et al.*, 1977) in hippocampus, while in invertebrates a classic model is presynaptic facilitation (Hawkins *et al.*, 1983; Heisenberg, 2003). Early-phase LTP and LTD are transient changes in synaptic strength that require the activation of protein kinases (Malenka *et al.*, 1989) and phosphatases (Mulkey *et al.*, 1993), while late-phase LTP and LTD are enduring changes that also require protein synthesis for their maintenance (Krug *et al.*, 1984; Manahan-Vaughan *et al.*, 2000; Roberson *et al.*, 1996). Equivalent early- and late-phase plasticity is also observed during presynaptic facilitation (Bailey *et al.*, 1992; Tully *et al.*, 1994; Hawkins *et al.*, 2006). A characteristic feature of memory is the so-called spacing effect, in which repeated stimuli must be spaced apart rather than massed together in time in order to promote the transition from short- to long-term memory (Ebbinghaus, 1885). Early- and late-phase synaptic plasticity are the physiological correlates of short- and long-term memory (Carew *et al.*, 1972; Tully *et al.*, 1994; Beck *et al.*, 2000; Sutton *et al.*, 2002), so a central challenge is to understand the molecular mechanisms by which spaced rather than massed stimulation protocols lead to late- rather than just early-phase plasticity. The unconditioned stimulus in a classical conditioning protocol activates neuromodulatory pathways in *Aplysia* (Hawkins *et al.*, 2006) and *Drosophila* (Heisenberg, 2003; Davis, 2005) and dopaminer-

gic activity induces protein synthesis-dependent potentiation in hippocampus, with strong, tetanising stimuli thought likely to co-activate dopaminergic pathways (Frey *et al.*, 1993; Huang & Kandel, 1995; Sajikumar & Frey, 2004). In all these cases, the enzyme adenylyl cyclase is a target for coincident neurotransmitter and neuromodulatory activity, producing the second messenger cAMP, and activating multiple signalling cascades including the CREB and MAPK pathways (Bartsch *et al.*, 1995; Yin *et al.*, 1995; Barco *et al.*, 2002). Spaced repetition generates repeated waves of MAPK activity, with modification of these waves via the phosphatase SHP2 modulating the interval required between repetitions (Pagani *et al.*, 2009).

In a mathematical model of integrative synaptic plasticity, we found evidence for the emergence of a natural timescale that hinted at a partial explanation of the spacing effect (Elliott & Lagogiannis, 2012). Such models propose that single synapses act as low-pass filters (Elliott, 2008), integrating plasticity induction signals and expressing synaptic plasticity only when the filter reaches an upper or lower threshold. These "integrate-and-express" models strongly control potentially destabilising fluctuations in synaptic strength in a developmental context (Elliott, 2008; Elliott & Lagogiannis, 2009). When applied to an associative memory task with discrete-state synapses with two or only a few discrete states of synaptic strength, for which experimental evidence is accumulating (Petersen *et al.*, 1998; O'Connor *et al.*, 2005; Montgomery & Madison, 2004; Bartol *et al.*, 2015), we found that they outperform non-integrative models in almost all biologically relevant regions of parameter space (Elliott & Lagogiannis, 2012; Elliott, 2016a; Elliott, 2016b). In all existing non-integrative models, the trace associated with a memory falls monotonically in time, dying way. Remarkably, in integrative models, the trace initially rises, reaches a peak determined by the filter size, and then finally dies away. The peak of this

trace is a natural time at which to repeat the presentation of the memory for re-storage, although neither neuromodulation nor the transition from early- to late-phase transition were considered (Elliott & Lagogiannis, 2012). However, we have proposed that an integrative synaptic filter could be instantiated on the relatively small number of macromolecules that are involved in plasticity at single synapse (Harris & Stevens, 1989; Nusser *et al.*, 1998; Bagal *et al.*, 2005; Miller *et al.*, 2005; Asrican *et al.*, 27), and specifically by collective protein kinase–phosphatase switches, such as the CaMKII-PP1 switch (Lisman & Zhabotinsky, 2001; Pi & Lisman, 2008), which permit decisive, all-or-none, switch-like responses to graded inputs (Ferrell, 1996; Bhalla & Iyengar, 1999; Alon, 2006). Such a kinase–phosphatase switch may therefore be one downstream target of the MAPK pathway that is activated by neuromodulatory signalling and is known to be involved in regulating the spacing effect (Pagani *et al.*, 2009). If our proposed integrative synaptic filter is under dynamic control via the neuromodulated balance between constitutive kinase and phosphatase activity, then dynamic integrative synaptic plasticity would provide an explanation for the spacing effect. In this context, we understand the spacing effect to reflect the capacity of single synapses to discriminate between massed and spaced repetitions of strong stimuli, a view supported by experimental evidence (Martin *et al.*, 2017).

Here, therefore, we modify our mathematical model of filter-based synaptic integration to include a regulated, dynamic decay process that is modulated by coincident neurotransmitter and neuromodulatory input during strong simulation protocols. With a very simple, abstract model that seeks to capture the possible role of the change in balance between kinase and phosphatase activity in controlling filter dynamics, we show that a synapse that dynamically integrates plasticity induction events can robustly distinguish between massed

and spaced repetition protocols, suppressing the response to massed protocols without significantly affecting the response to spaced protocols. We examine the parameter space of the model in detail, establishing the values of parameters and spaced repetition times that optimise the ability of a single synapse to discriminate between spaced and massed repetition protocols. Given our assumption that the spacing effect reflects the capacity of single synapses to robustly discriminate between spaced and massed repetition protocols, these optimally spaced repetition times in our model represent concrete experimental predictions for stimulus repetition times that maximally promote the transition from early- to late-phase plasticity and thus the formation of long-term memory. We show that the model's response to spaced repetitions is robust to changes around these optimal parameters and repetition times. However, under large changes in its parameters, the model can enhance the response to massed protocols or suppress the response to any protocol. The model therefore predicts either that massed protocols may also promote long-term memory formation or that spaced protocols promote only short-term memory formation under pharmacological manipulations that emulate these model parameter changes, consistent with experimental results (Pagani *et al.*, 2009).

2 Background and Setup

2.1 Perceptron Formulation

We consider a single perceptron with N synapses of strengths $S_i(t)$, $i = 1, \ldots, N$, at time $t \geq 0$ s taking discrete values from the set $\{\Omega_1, \ldots, \Omega_\nu\}$, where $\Omega_a = -1 + 2(a-1)/(\nu-1)$. We initially consider $\nu \geq 2$ but will restrict to binary-strength synapses with $\nu = 2$. This arrangement of strengths is conventional and does not imply that negative (positive) values correspond

to inhibitory (excitatory) synapses because we may scale them at will provided we modify the perceptron's firing threshold ϑ appropriately. We may assume that the perceptron experiences binary-valued inputs $x_i \in \{-1, +1\}$ through these N synapses, with the input vector $\boldsymbol{x}$ generating the activation or unthresholded output

$$h_{\boldsymbol{x}}(t) = \frac{1}{N} \sum_{i=1}^{N} x_i \, S_i(t) \tag{1}$$

as standard (Hertz *et al.*, 1991). Again, we may scale these input values provided we adjust ϑ accordingly. As we are only concerned with whether or not $h_{\boldsymbol{x}}(t)$ exceeds ϑ, we need not consider any non-linearity that converts $h_{\boldsymbol{x}}(t)$ into an explicit firing rate.

The perceptron is required to store sequentially the "memories" $\boldsymbol{\xi}^{\alpha}$, $\alpha = 0, 1, 2, \ldots$, at times governed by a Poisson process of rate r Hz. A continuous-time process is more realistic than a discrete-time process, and the Poisson process is the easiest to examine. Input ξ_i^{α} is the plasticity induction signal to synapse i upon storage of memory α. These induction signals may induce either plastic (expressed) or metaplastic (unexpressed) changes at a synapse, with $\xi_i^{\alpha} = +1$ ($\xi_i^{\alpha} = -1$) corresponding to a potentiating (depressing) induction signal for firing threshold $\vartheta > 0$. The signals ξ_i^{α} are as usual taken to be independent across synapses $i = 1, \ldots, N$ and between memories $\alpha = 0, 1, 2, \ldots$ (Hertz *et al.*, 1991), with $\mathsf{Prob}[\xi_i^{\alpha} = \pm 1] = g_{\pm}$. So that potentiation and depression processes are equiprobable and treated completely symmetrically, we take $g_{\pm} = \frac{1}{2}$. Memory $\boldsymbol{\xi}^0$ is always stored at time $t = 0$ s with all later memories stored at times $t > 0$ s. The storage of a memory may cause changes in synaptic strengths, so the storage of later memories may degrade the recall of earlier memories. We are thus concerned specifically with the fidelity of recall

of the initial memory $\boldsymbol{\xi}^0$ in the face of this ongoing storage of later memories, and we refer to $\boldsymbol{\xi}^0$ as the tracked memory. We write $h(t) = h_{\boldsymbol{\xi}^0}(t)$ for simplicity, where this tracked memory signal is the perceptron's activation in response to the re-presentation and not the re-storage of $\boldsymbol{\xi}^0$ at some later time $t > 0$ s.

2.2 Integrate-and-Express Synapses

In an integrate-and-express model of synaptic plasticity, the induction signals at each synapse do not necessarily lead to the immediate expression of plasticity. Rather, each synapse instantiates a filter that integrates its induction signals, with the expression of synaptic plasticity occurring only when the filter reaches an upper or lower threshold. As we consider $g_\pm = \frac{1}{2}$, we also consider symmetric filters in which the upper and lower thresholds are $\pm\Theta$, respectively. Labelling the state of a synapse's filter with letters such as I and J, with $I, J \in \{-(\Theta-1), \ldots, +(\Theta-1)\}$, we have the transitions

$$\begin{aligned} \xi_i^\alpha = +1 \Rightarrow \begin{cases} I \mapsto I+1 & \text{for } I < +(\Theta-1) \\ I \mapsto 0 \;\&\; \Uparrow & \text{for } I = +(\Theta-1) \end{cases} \\ \xi_i^\alpha = -1 \Rightarrow \begin{cases} I \mapsto I-1 & \text{for } I > -(\Theta-1) \\ I \mapsto 0 \;\&\; \Downarrow & \text{for } I = -(\Theta-1) \end{cases} \end{aligned} \quad , \tag{2}$$

where potentiating (depressing) induction signals increment (decrement) the filter state, but if the filter thereby reaches its upper (lower) threshold, then its state is reset to zero and potentiation (depression) is expressed if possible, indicated by the "$\Uparrow$" ("$\Downarrow$") symbol. We could consider alternative resetting dynamics (Elliott, 2016b), but resetting to zero is the most natural for synapses with general $\nu \geq 2$. We may represent the state of a synapse's filter by a $(2\,\Theta - 1)$-dimensional vector whose components are (unconventionally) indexed by I.

Changes in filter state may then be implemented by the $(2\,\Theta-1)\times(2\,\Theta-1)$ matrices $\mathbb{F}^{\pm}$ (for non-threshold processes) and $\mathbb{T}^{\pm}$ (for threshold processes) for induction signals $\xi_i^\alpha = \pm 1$, respectively. The matrix $\mathbb{F}^{+}$ ($\mathbb{F}^{-}$) contains elements of unity on its lower (upper) diagonal and zeros elsewhere, and just steps the filter state up (down) without resetting to zero. The two matrices $\mathbb{T}^{\pm}$ have elements $\mathbb{T}^{\pm}_{0,\pm(\Theta-1)} = 1$ (correlated signs) with all others being zero, and reset the filter state to zero from upper or lower threshold. Elements of filter matrices are again (also unconventionally) indexed by I and J.

Potentiation (depression) is expressed by synapse i's strength increasing (decreasing) by $2/(\nu-1)$ when its filter reaches upper (lower) threshold. However, if $S_i = +1$ (-1) with $\xi_i^\alpha = +1$ (-1), then the synapse is already saturated at its upper (lower) strength, and no further change is possible. As there are ν possible strength states, we may represent the joint strength and filter state of a synapse by a $(2\,\Theta-1)\nu$-dimensional vector, say $\boldsymbol{P}(t)$, with changes in this joint synaptic state implemented by $(2\,\Theta-1)\nu\times(2\,\Theta-1)\nu$ transition matrices, say $\mathbb{M}^{\pm}$ for $\xi_i^\alpha = \pm 1$, respectively. Letting the a^{th} block of $2\,\Theta-1$ entries of $\boldsymbol{P}(t)$ represent the filter state when the synapse has strength Ω_a, $a = 1,\ldots,\nu$, the matrices $\mathbb{M}^{\pm}$ may be written schematically in block forms as

$$\mathbb{M}^{+} = \left(\begin{array}{c|c|c|c|c} \mathbb{F}^{+} & & & & \\ \hline \mathbb{T}^{+} & \mathbb{F}^{+} & & & \\ \hline & \ddots & \ddots & & \\ \hline & & \mathbb{T}^{+} & \mathbb{F}^{+} & \\ \hline & & & \mathbb{T}^{+} & \mathbb{F}^{+}+\mathbb{T}^{+} \end{array}\right) \text{ and } \mathbb{M}^{-} = \left(\begin{array}{c|c|c|c|c} \mathbb{T}^{-}+\mathbb{F}^{-} & \mathbb{T}^{-} & & & \\ \hline & \mathbb{F}^{-} & \mathbb{T}^{-} & & \\ \hline & & \ddots & \ddots & \\ \hline & & & \mathbb{F}^{-} & \mathbb{T}^{-} \\ \hline & & & & \mathbb{F}^{-} \end{array}\right), \tag{3}$$

where the elements are zero unless explicitly indicated. These two matrices implement all possible, allowed, simultaneous transitions in joint strength and filter state at a synapse experiencing induction signal $\xi_i^\alpha = \pm 1$.

If one element of $\boldsymbol{P}(t)$ is unity and all others are zero, corresponding to a definite strength and filter state at time t, then either $\mathbb{M}^+\boldsymbol{P}(t)$ or $\mathbb{M}^-\boldsymbol{P}(t)$ also corresponds to a definite synaptic state for a synapse i experiencing a definite induction signal of either $\xi_i^\alpha = +1$ or $\xi_i^\alpha = -1$, respectively, at time t. However, memories are presented at Poisson times. Further, we are not interested in any particular tracked memory $\boldsymbol{\xi}^0$ nor in any particular set of non-tracked memories $\boldsymbol{\xi}^\alpha$, $\alpha \geq 1$. Rather, we are interested only in the typical behaviour of $h(t)$ for a typical tracked memory and typical non-tracked memories. Just as the memory presentation times are random or stochastic variables, the components ξ_i^α are therefore also stochastic variables. Hence, even if $\boldsymbol{P}(0)$ is definite, $\boldsymbol{P}(t)$ for $t > 0$ s corresponds to a joint probability distribution for a synapse's later joint strength and filter state. The tracked memory signal $h(t)$ is therefore also a stochastic variable, and rather than its behaviour for any particular realisation of the memories $\boldsymbol{\xi}^\alpha$ and their Poisson times, we are interested in its statistical properties averaged over all possible memories. We are especially concerned with its mean, $\mu(t) = \mathsf{E}[h(t)]$, and to a lesser extent its variance, $\sigma^2(t) = \mathsf{Var}[h(t)]$.

To average over all later memories with $\alpha \geq 1$, we do not consider a definite induction signal ξ_i^α but instead the probability distribution of these signals. Hence, for memory $\alpha \geq 1$, synapse i experiences the superposed transition matrix

$$\mathbb{M} = g_+\mathbb{M}^+ + g_-\mathbb{M}^-, \tag{4}$$

so that the effects of both $\xi_i^\alpha = +1$ and $\xi_i^\alpha = -1$ are considered, weighted by their probabilities, here $g_\pm = \frac{1}{2}$. For the initial memory $\boldsymbol{\xi}^0$, synapse i may experience at $t = 0$ s either $\xi_i^0 = +1$ or $\xi_i^0 = -1$. If the state of the synapse immediately prior to the presentation of $\boldsymbol{\xi}^0$ is $\boldsymbol{A}$, then its state $\boldsymbol{P}(0)$

immediately after is a mixture (not a superposition) of the two states $\mathbb{M}^+\boldsymbol{A}$ and $\mathbb{M}^-\boldsymbol{A}$. The state $\boldsymbol{A}$ is assumed to be governed by the equilibrium distribution of a synapse's joint strength and filter state. It is therefore the eigenvector of the matrix $\mathbb{M}$ with unit eigenvalue, normalised to a probability distribution. From the block structure of $\mathbb{M}$, this eigenvector is determined by the unit eigenvector of the combined submatrix $\frac{1}{2}\,(\mathbb{T}^- + \mathbb{F}^- + \mathbb{F}^+ + \mathbb{T}^+)$. Denoting this latter eigenvector by $\boldsymbol{\mathcal{A}}$, it has normalised components $\mathcal{A}_I = (\Theta - |I|)/\Theta^2$, and so $\boldsymbol{A}$ is given in schematic block form by

$$\boldsymbol{A} = \frac{1}{\nu}\big(\underbrace{\boldsymbol{\mathcal{A}}^{\mathrm{T}}|\ldots|\boldsymbol{\mathcal{A}}^{\mathrm{T}}}_{\nu}\big)^{\mathrm{T}}, \tag{5}$$

where the superscript T denotes the transpose. Thus, the distribution of filter states is independent of strength state in equilibrium. Defining the matrix $\mathbb{G} = \mathbb{M} - \mathbb{I}$, where $\mathbb{I}$ is the identity matrix, the distribution of a synapse's joint strength and filter state evolves according to the Master equation

$$\frac{d\boldsymbol{P}(t)}{dt} = r\,\mathbb{G}\boldsymbol{P}(t), \tag{6}$$

with solution $\boldsymbol{P}(t) = e^{rt\mathbb{G}}\boldsymbol{P}(0)$, where $\boldsymbol{P}(0)$ is either $\mathbb{M}^+\boldsymbol{A}$ or $\mathbb{M}^-\boldsymbol{A}$, depending on $\xi_i^0 = \pm 1$.

To determine the statistics of $h(t)$ we require the two conditional probabilities $\mathsf{Prob}[S_i(t) = \Omega_a \,|\, \xi_i^0 = \pm 1]$, $a = 1, \ldots, \nu$, which can be obtained at once from $\boldsymbol{P}(t)$ by summing over the a^{th} block of its components for the appropriate choice of $\boldsymbol{P}(0) = \mathbb{M}^{\pm}\boldsymbol{A}$ for $\xi_i^0 = \pm 1$. When potentiation and depression processes are treated completely symmetrically and the possible strengths $\Omega_1, \ldots, \Omega_\nu$ are symmetrically arranged around zero, we can show that while the two distributions $\mathsf{Prob}[S_i(t) = \Omega_a \,|\, \xi_i^0 = \pm 1]$ depend on ξ_i^0, the two distributions

$\mathsf{Prob}[\xi_i^0 S_i(t) = \Omega_a \,|\, \xi_i^0 = \pm 1]$ are identical and independent of the sign of ξ_i^0 (Elliott, 2016a). Since the tracked memory signal $h(t) = h_{\boldsymbol{\xi}^0}(t) = \frac{1}{N}\sum_{i=0}^{N} \xi_i^0 S_i(t)$, it is therefore the average of N identically distributed random variables, regardless of the tracked memory $\boldsymbol{\xi}^0$. Thus, in terms of their contributions to $h(t)$, the two parts in the equiprobable mixture of $e^{rt\mathbb{G}}\mathbb{M}^{\pm}\boldsymbol{A}$ contribute identically. This means that for determining the statistics of $h(t)$, we need only consider a definite choice of either $\mathbb{M}^{+}\boldsymbol{A}$ or $\mathbb{M}^{-}\boldsymbol{A}$ and not, in fact, a mixture of them. These observations significantly simplify calculations. We always choose $\mathbb{M}^{+}\boldsymbol{A}$. Because $h(t)$ is the average of N identically distributed random variables, its mean is just $\mu(t) = \mathsf{E}[\xi_i^0 S_i(t)]$ for any single synapse i, with $\mu(t)$ being independent of N. Further, with the choice of $\mathbb{M}^{+}\boldsymbol{A}$ from the initial mixture, we may just set $\xi_i^0 = +1$, so that $\mu(t) = \mathsf{E}[S_i(t)]$ for any i.

With the problem so defined, we may explicitly compute the average perceptron activation $\mu(t)$ in response to re-presentation of the tracked memory at some future time $t > 0$ s. Extensive calculations performed elsewhere (Elliott, 2016a) produce the result,

$$\mu(t) \;=\; \frac{4}{\Theta^3\,\nu(\nu-1)}\left(\frac{1}{\nu}\sum_{l=0}^{\left\lfloor\frac{\Theta\,\nu-1}{2}\right\rfloor}\cot^2\tfrac{(2\,l+1)\pi}{2\,\Theta\,\nu}\,\exp\left\{-rt\left[1-\cos\tfrac{(2\,l+1)\pi}{\Theta\,\nu}\right]\right\}\right.$$
$$\left. -\,\nu\sum_{l=0}^{\left\lfloor\frac{\Theta-1}{2}\right\rfloor}\cot^2\tfrac{(2\,l+1)\pi}{2\,\Theta}\,\exp\left\{-rt\left[1-\cos\tfrac{(2\,l+1)\pi}{\Theta}\right]\right\}\right), \tag{7}$$

where $\lfloor\,\cdot\,\rfloor$ denotes the floor function. In radical contrast to all existing non-integrative memories of synaptic plasticity applied to memory storage, rather than monotonically decreasing with time, $\mu(t)$ initially increases, reaches a

maximum at a time given approximately by

$$rt_{\text{peak}} \approx \frac{4\,\Theta^2}{\pi^2}\frac{\nu^2}{\nu^2-1}\log_e \nu, \tag{8}$$

and then decays to zero (its equilibrium value) in a strongly ν-dependent manner, exhibiting near plateau-like behaviour for larger values of ν before decaying (Elliott, 2016a). The initial increase in $\mu(t)$ is actually facilitated or driven by the ongoing storage of later memories, while yet later memory storage eventually erodes the recall of $\boldsymbol{\xi}^0$. In all other, non-integrative models, ongoing memory storage always erodes the recall of $\boldsymbol{\xi}^0$.

2.3 Dynamic Synaptic Integration

The model of synaptic integration discussed above is non-dynamical in the sense that in the absence of further input, a synapse's filter state would remain fixed and unchanged. However, we might naturally expect it to decay back to the zero state in the absence of activity. For example, as discussed, we can imagine that the filter state is instantiated by the phosphorylation state of a small collection of macromolecules at single synapses. If the natural, basal state of one of these molecules is to be singly phosphorylated, then a potentiating induction signal may lead to its becoming doubly phosphorylated while a depressing induction signal may lead to its becoming dephosphorylated. In addition to this activity-dependent phosphatase and kinase activity, constitutively active phosphatases and kinases may restore such a molecule on average to the singly phosphorylated, basal state. Thus, for one such molecule, say $\mathcal{M}$,

we can imagine the following abstract reaction scheme,

$$\mathcal{M} \underset{\xrightarrow{rg_+}}{\overset{\xleftarrow{rg_-}}{\xrightarrow{\eta}}} \mathcal{M}^* \underset{\xrightarrow{rg_+}}{\overset{\xleftarrow{rg_-}}{\xleftarrow{\eta}}} \mathcal{M}^{**} \tag{9}$$

in which $\mathcal{M}$, $\mathcal{M}^*$ and $\mathcal{M}^{**}$ represent the dephosphorylated, singly phosphorylated and doubly phosphorylated states, respectively; η denotes the constitutive decay rate back to the basal state; and $rg_\pm$ denote the rates of potentiating and depressing induction signals. A filter state $I > 0$ would correspond to $|I|$ of these molecules being doubly phosphorylated while $I < 0$ would correspond $|I|$ of them being dephosphorylated. We would then expect filter state $\pm I$ $(I > 0)$ to decay to state $\pm(I-1)$ at a rate governed by $I\eta$, since any one of the I dephosphorylated or doubly phosphorylated molecules reverts to its singly phosphorylated state at rate η.

Let $\mathbb{X}$ denote the $(2\,\Theta-1)\times(2\,\Theta-1)$ generating matrix that implements this decay process for filter states, without the overall rate factor η. For example, for $\Theta = 3$, $\mathbb{X}$ is given by

$$\mathbb{X} = \begin{pmatrix} -2 & 0 & 0 & 0 & 0 \\ +2 & -1 & 0 & 0 & 0 \\ 0 & +1 & 0 & +1 & 0 \\ 0 & 0 & 0 & -1 & +2 \\ 0 & 0 & 0 & 0 & -2 \end{pmatrix}. \tag{10}$$

We then write down the $(2\,\Theta-1)\nu\times(2\,\Theta-1)\nu$ decay matrix $\mathbb{D}$ in the schematic block form

$$\mathbb{D} = \left(\begin{array}{c|c|c|c|c} \mathbb{X} & & & & \\ \hline & \mathbb{X} & & & \\ \hline & & \ddots & & \\ \hline & & & \mathbb{X} & \\ \hline & & & & \mathbb{X} \end{array}\right), \tag{11}$$

where this matrix implements decay of a synapse's filter state regardless of its strength state. Then, with both ongoing memory storage at a fixed rate r and filter state decay occurring at a possibly time-dependent rate $\eta(t)$, a synapse's joint strength and filter state now evolves according to the Master equation

$$\frac{d\boldsymbol{P}(t)}{dt} = \big[r\,\mathbb{G} + \eta(t)\,\mathbb{D}\big]\boldsymbol{P}(t). \tag{12}$$

The transitions encoded by the matrix $r\,\mathbb{G} + \eta(t)\,\mathbb{D}$ in this equation are represented graphically in Fig. 1. If $\eta(t)$ is in fact a constant, independent of time, then we still have the standard exponential solution of the form $\boldsymbol{P}(t) = e^{t(r\,\mathbb{G}+\eta\,\mathbb{D})}\boldsymbol{P}(0)$. But for time-dependent $\eta(t)$, because the two matrices $\mathbb{G}$ and $\mathbb{D}$ do not commute, we require the time-ordered exponential solution. For our purposes here such a formal solution is essentially useless and we must resort to numerical methods to solve Eq. (12).

Physiologically, potentiating and depressing stimuli appear to be divided into two classes: strong stimuli that simultaneously activate both neurotransmitter and neuromodulatory systems, and weak stimuli that activate neurotransmitter systems without also activating neuromodulatory systems (Frey *et al.*, 1993; Huang & Kandel, 1995; Roberson *et al.*, 1996; Heisenberg, 2003; Davis, 2005; Hawkins *et al.*, 2006). Strong, repetitive stimuli that are spaced apart are required for late-phase plasticity and the conversion of short- into long-term memory, while either strong, repetitive stimuli that are massed together or weak stimuli lead only to early-phase plasticity and short-term memory that does not endure (Carew *et al.*, 1972; Tully *et al.*, 1994; Beck *et al.*, 2000; Sutton *et al.*, 2002). Accordingly, we model the tracked memory $\boldsymbol{\xi}^0$ as a strong stimulus, perhaps one of particular salience or importance that also activates neuromodulatory systems. Memory $\boldsymbol{\xi}^0$ is taken as a strong stimu-

lus whether stored initially at time $t = 0$ s or repeated and re-presented for re-storage at later times. Conversely, the later memories $\boldsymbol{\xi}^{\alpha}$ for $\alpha \geq 1$ are modelled as weak stimuli that do not co-activate neuromodulatory systems. We may now also think of these later memories as reflecting the (perhaps random) background or spontaneous synaptic activity arising at the Poisson rate r on which is superimposed the strong stimuli corresponding to the storage of $\boldsymbol{\xi}^{0}$.

We model the dynamic decay rate $\eta(t)$ as satisfying the simple equation

$$\frac{d\eta(t)}{dt} = -\frac{1}{\tau}\left[\eta(t) - \overline{\eta}\right], \tag{13}$$

where τ^{-1} is an inverse timescale that determines the rate at which $\eta(t)$ returns to its equilibrium or baseline value, $\overline{\eta}$. We could take $\overline{\eta} > 0$ Hz, but this would only complicate yet not fundamentally modify our analysis. Instead, we set $\overline{\eta} = 0$ Hz for simplicity and convenience. We further suppose that in the presence of a strong stimulus, $\eta(t)$ undergoes an increase that is rapid on a timescale compared to τ. Thus, a strong stimulus dynamically adjusts the constitutive filter decay rate, making the filter state decay more quickly for a period of time determined by both the timescale τ and the size of the rapid increase in $\eta(t)$. For the strong stimulus $\boldsymbol{\xi}^{0}$ presented or re-presented at time t, we write

$$\eta(t) \mapsto \eta(t) + \Delta\eta, \tag{14a}$$

$$\boldsymbol{P}(t) \mapsto \mathbb{M}^{+}\boldsymbol{P}(t), \tag{14b}$$

so that $\eta(t)$ experiences an instantaneous jump by an amount $\Delta\eta$, and of course the joint distribution of a strongly stimulated synapse's strength and filter state also undergoes a change in response to the induction stimulus, described by

the second of these two equations. These two equations also hold at $t = 0$ s, with the second being equivalent to $\boldsymbol{P}(0) = \mathbb{M}^{+}\boldsymbol{A}$, which represents the initial storage of $\boldsymbol{\xi}^0$. If memory $\boldsymbol{\xi}^0$ is stored at times $t = t_0, \ldots, t_\rho$ where $\rho \geq 0$ is the number of times that the memory is repeated for re-storage, and we always take $t_0 = 0$ s, then $\eta(t)$ is given by

$$\eta(t) = \left[\sum_{i=0}^{\rho} H(t - t_i)e^{-(t-t_i)/\tau}\right]\Delta\eta, \tag{15}$$

where $H(t)$ is the Heaviside step function, and to avoid ambiguity we define $H(0) \equiv 1$. The advantage of the simple form in Eq. (13) and the impulse-like character of Eq. (14a) is that they admit of the simple solution in Eq. (15). We would not expect any significant modifications to our results if we modelled the rapid rise in $\eta(t)$ with each strong stimulus as for example an α-function rather than an impulse, so we use impulses for convenience.

For a single strong stimulus applied only at $t = 0$ s, we note that $\int_0^\infty dt\, \eta(t) = \tau\Delta\eta$, which depends on τ. If $\eta(t)$ represents some measure of the total production or availability of a protein kinase or phosphatase as a result of the strong stimulus, then it may be argued that we should instead have, say, $\int_0^\infty dt\, \eta(t) = \eta_0$, where η_0 is both independent of τ and a dimensionless quantity. For this, we would need to take $\eta(t) = \frac{\eta_0}{\tau}e^{-t/\tau}$ (for a strong stimulus at $t_0 = 0$ s), where the function $\frac{1}{\tau}e^{-t/\tau}$ has unit integral over $t \geq 0$. Such a form would facilitate direct comparison to alternative profiles such as an α-function since in this case the integral would be a constant η_0, independent of profile. Each normalisation convention has its merits, but for an impulse it suffices to define $\Delta\eta = \eta_0/\tau$. We use both $\Delta\eta$ and η_0 to illustrate different properties.

Because $\overline{\eta} = 0$ Hz, in the absence of further strong stimulation $\eta(t)$ decays to zero, so that for large times, the filter decay process drops out and we are

left with dynamics governed only by the ongoing storage of the weak memories $\boldsymbol{\xi}^\alpha$, $\alpha \geq 1$. The equilibrium distribution $\boldsymbol{A}$ of a synapse's strength and filter states is therefore not affected by the filter decay process, so we continue to use $\boldsymbol{A}$ in Eq. (5) as the background against which the tracked, strong memory $\boldsymbol{\xi}^0$ is stored initially at time $t_0 = 0$ s. Of course, for all later repetitions of $\boldsymbol{\xi}^0$ at times t_i, $i > 0$, the distribution $\boldsymbol{P}(t_i)$ changes in accordance with Eq. (14b).

To reduce the free parameters, we write $t' = rt$, absorbing the rate r into a re-definition of time. We must also write $t'_i = rt_i$ for the times of the strong stimuli $\boldsymbol{\xi}^0$ and also define $\tau' = r\tau$, as all timescales must be so re-defined. Because η itself is a rate, we further use r to define $\eta' = \eta/r$, so that η' is dimensionless. Then $\eta' \mapsto \eta' + \Delta\eta'$, where $\Delta\eta' = (\Delta\eta)/r$ is now the size of the impulse applied to η' during strong stimulation. Alternatively, because $\Delta\eta = \eta_0/\tau$, we also have $\Delta\eta' = \eta_0/\tau'$, so that the re-definition of temporal quantities leaves the dimensionless η_0 invariant, as it should. With these conventions, our system of equations becomes

$$\frac{d\boldsymbol{P}(t')}{dt'} = [\mathbb{G} + \eta'(t')\mathbb{D}]\,\boldsymbol{P}(t'), \tag{16a}$$

$$\frac{d\eta'(t')}{dt'} = -\frac{1}{\tau'}\,\eta'(t'), \tag{16b}$$

$$\boldsymbol{P}(t'_i) \mapsto \mathbb{M}^+\boldsymbol{P}(t'_i), \tag{16c}$$

$$\eta'(t'_i) \mapsto \eta'(t'_i) + \Delta\eta' \text{ or } \eta'(t'_i) + \eta_0/\tau', \tag{16d}$$

where the first two equations describe the generic evolution of $\boldsymbol{P}(t')$ and $\eta'(t')$ at all times $t \geq 0$, and the last two describe the instantaneous changes in them at the times $t'_0, \ldots, t'_\rho$ of the storage of the strong memory $\boldsymbol{\xi}^0$. We also work with the tracked memory signal $h(t')$ and its mean $\mu(t')$, etc. The result of these considerations is that we may show results for $\mu(rt)$ and $\eta(rt)/r$ plotted against

rt, with decay model parameters specified by $r\tau$ and impulses for $\eta(rt)/r$ of either $\Delta\eta/r$ or $\eta_0/(r\tau)$. In this way, we essentially eliminate the rate r since all other quantities scale with it in the manner indicated. We may then just set $r = 1$ Hz without loss of generality in solving the system in Eq. (16), and then reinstate r as described.

Although $\eta(t)$ has been introduced in a somewhat abstract way, we may think of it in concrete terms by supposing that its level in excess of baseline represents the level of activity in the downstream signalling cascades activated by the second messenger cAMP upon coincident neurotransmitter and neuromodulatory signalling. Thinking specifically of the MAPK pathway, spaced repetition generates waves of activity in this pathway (Pagani *et al.*, 2009). Therefore, we may think of the two parameters $\Delta\eta$ and τ as controlling the amplitude and time-course of these waves, respectively, at least in abstract terms. The filter decay rate $\eta(t)$ may then be thought of as some level of kinase and phosphatase activity in excess of baseline, perhaps identified in particular with the phosphatase SHP2, which is known to regulate the spacing effect (Pagani *et al.*, 2009). We stress that we need not be committed to this particular view, but it provides a way of thinking about our abstract model in concrete, mechanistic, underlying terms.

3 Results

We have formulated the model with arbitrary $\nu \geq 2$ for generality, but we now restrict to the specific case of binary-strength synapses with $\nu = 2$. Synapses with strength $S_i = -1$ may be thought of as "weak" synapses and those with $S_i = +1$ as "strong" synapses; as discussed, negative values do not imply inhibition but merely reflect a mathematically convenient convention after rescaling.

We restrict to the $\nu = 2$ case because the fundamental filter dynamics are amply exhibited by binary-strength synapses. As can be seen from Eq. (8), for example, the principal behaviour is captured by the strong (there quadratic) dependence on filter size Θ, while the dependence on the number of states of strength ν is very mild (there logarithmic). For the purposes of discussing our results, considering $\nu > 2$ adds nothing fundamentally new compared to the particular case of $\nu = 2$.

3.1 Massed Versus Spaced Repetition Without Decay

We first consider memory dynamics in which the filter decay mechanism is suppressed, setting $\eta'(t') \equiv 0$ or equivalently $\Delta\eta' = 0$ so that the filter decay rate does not receive an impulse with each strong stimulus. For massed repetition, we store memory $\boldsymbol{\xi}^0$ regularly and uniformly at the early times $t'_i = i$ for $i = 0, \ldots, \rho$, so a total of $1 + \rho$ times including at $t' = 0$ and repeated a further ρ times. For spaced repetition, we use the successive mean memory signal peaks to define t'_i. Presenting just once at $t'_0 = 0$, we induce a first peak at t'_{peak} given by Eq. (8). This defines t'_1 as the time at which $\boldsymbol{\xi}^0$ is stored for a second time. A given set of times $\{t'_0, \ldots, t'_i\}$ therefore automatically induces a later repetition time t'_{i+1}. Eq. (8) applies only to t'_1, so all repetition times are determined numerically. We also consider for reference a random repetition scenario in which the times $t'_1, \ldots, t'_\rho$ are generated randomly, with each inter-repetition interval drawn from a uniform distribution between unity and twice the average at-peak repetition interval.

Typical results are shown in Fig. 2. In panels A, C and E we show the dynamics of $\mu(t')$ for the representative choice of $\Theta = 8$, using $\rho = 12$ repetitions of $\boldsymbol{\xi}^0$. The overall envelope in each case shows the dynamics of $\mu(t')$ in response to all repetitions, while the various curves below this overall envelope

show what the dynamics would have been had there been fewer repetitions. For massed repetitions (panel A), there is a very large and very rapid rise in $\mu(t')$. For at-peak repetitions (panel B), $\mu(t')$ progressively rises with each repetition. The maximum of the overall envelope of course occurs much later than that for massed repetitions, but the at-peak maximum is lower than that for massed repetitions. For random repetitions (panel C), $\mu(t')$ can go down as well as up between repetitions, depending on whether or not a repetition catches $\mu(t')$ before or after a peak, but its overall behaviour is somewhat similar to that for at-peak repetitions, with the overall locations and amplitudes of the maxima being similar, at least for this particular set of random repetition times. Panels B, D and F compare massed and at-peak repetitions for different filter sizes and different numbers of repetitions by showing the maximum value of $\mu(t')$ attained during a particular stimulation protocol. This maximum necessarily saturates as the number of repetitions increases, but for larger filters, the dependence of this maximum on the number of repetitions is linear or nearly so. For smaller filters, the saturation dynamics are very marked, especially for massed repetitions. Comparing massed and at-peak protocols directly in panel F, massed repetition *always* induces a larger maximum of $\mu(t')$ than at-peak repetition, for any choice of Θ and any choice of $\rho \geq 1$. Without a dynamic decay mechanism, we conclude that massed repetition is always more efficacious than spaced repetition in generating a very large perceptron activation, which translates into a higher firing rate in, say, a recognition task. The maximum for massed protocols always occurs much earlier than for spaced protocols, and so in principle massed protocols can always be distinguished from spaced protocols by a suitable temporal filtering window on the activation dynamics. However, a more simple threshold on the perception's firing rate could not distinguish between massed and spaced protocols and indeed may favour massed

over spaced protocols.

To understand these differences between massed and at-peak protocols in the absence of a filter decay mechanism, we examine in Fig. 3 the probability distribution of the filter states, conditioned on synaptic strength, during these protocols. We select $\Theta = 8$ again and $\rho = 6$ as representative. Panel A shows the probability distribution for a synapse to be in filter state I, given that the synapse is either weak ($S_i = -1$) or strong ($S_i = +1$), for both massed and at-peak protocols. These distributions are shown for different time points, corresponding to the seven times at which the memory $\boldsymbol{\xi}^0$ is stored and then times corresponding to overall signal peak and twice that time. Immediately prior to $t' = 0$, the distributions for weak and strong synapses are completely symmetric, and they are identical. The transition matrix $\mathbb{M}^+$, used for storing $\boldsymbol{\xi}^0$, then shifts these distributions upwards by one filter state. Weak synapses that reach the upper filter threshold are potentiated, becoming strong synapses, and their filter states are returned to zero; strong synapses reaching threshold cannot be potentiated further, but their filter states are also reset. This overall process occurs for each repetition of $\boldsymbol{\xi}^0$. (Note that had we instead chosen $\xi_i^0 = -1$ and $\mathbb{M}^-$, the argument would be precisely reversed, with downwards filter steps and strong synapses becoming weak.) For massed repetition, this repetitive upwards shifting of filter states occurs so quickly that relatively few other memories (with $\alpha \geq 1$) are also stored with ongoing memory storage. These other memories may have either $\xi_i^\alpha = +1$ or $\xi_i^\alpha = -1$, so on average their storage tends to pull the filter distribution in both directions, upwards and downwards, partially re-symmetrising the distribution, for both weak and strong synapses. With few such re-symmetrisation processes occurring during massed repetitions of $\boldsymbol{\xi}^0$, the repetitive influence of $\mathbb{M}^+$ rapidly leads to the potentiation of weak synapses and to the development of a large bias in filters

states, in both weak and strong synapses, towards the upper filter threshold. Hence, for massed repetition, we see a rapid, relentless, almost deterministic process of successively greater biasing with more repetitions (until the system essentially becomes saturated at near-maximal bias). In contrast, for at-peak repetition protocols and indeed spaced repetition protocols in general, although each repetition of $\boldsymbol{\xi}^0$ steps filter states upwards as described, the time between each repetition of $\boldsymbol{\xi}^0$ permits more than just a few of the other memories to be stored between the $\boldsymbol{\xi}^0$ stimuli. Hence, between each $\boldsymbol{\xi}^0$ stimulus, a significant degree of re-symmetrisation of filter states can occur, for both weak and strong synapses. Although repetition of $\boldsymbol{\xi}^0$ still leads to an overall increase in $\mu(t)$ and the development of upwards biasing in the filter states for both weak and strong synapses, this biasing is not as extensive as for massed protocols. We see these differences clearly in panel A of Fig. 3, up to the sixth repetition. Synapses experiencing a massed protocol have a very low probability of being in filter states with $I < 0$, while this is not true for a spaced protocol.

In panels B and C of Fig. 3, we quantify these observations by plotting $\mathsf{Prob}[I = \pm(\Theta - 1) \,|\, S_i = \mp 1]$ (in B) and the average "splitting" probability $\langle \pi_J^+ \rangle_J$ for weak and strong synapses. $\mathsf{Prob}[I = -(\Theta - 1) \,|\, S_i = +1]$ measures how likely a strong synapse is to be in the lowest filter state, while $\mathsf{Prob}[I = +(\Theta - 1) \,|\, S_i = -1]$ measures how likely a weak synapse is to be in the highest filter state. The splitting probabilities $\langle \pi_J^+ \rangle_J$ for $S_i = \pm 1$ measure how likely a weak or strong synapse is to escape through the upper filter threshold at some future time, averaged over all filter states, assuming only ongoing non-tracked memory storage. For massed protocols, there is a large difference between the two probabilities $\mathsf{Prob}[I = \pm(\Theta - 1) \,|\, S_i = \mp 1]$, with strong synapses having a very low probability of being in the lowest filter state compared to weak synapses being in the highest filter state. For spaced protocols, $\mathsf{Prob}[I =$

$-(\Theta-1)\,|\,S_i=+1]$ is nearly always non-negligible, except for immediately after the storage of $\boldsymbol{\xi}^0$. In panel C, the two splitting probabilities for the massed protocol are very clearly different, while for the spaced protocol they are extremely similar, again indicating large differences in biasing induced by these two protocols. We note that in panel B, for either protocol, when the two probabilities $\mathsf{Prob}[I=\pm(\Theta-1)\,|\,S_i=\mp 1]$ are equal following the final repetition of $\boldsymbol{\xi}^0$, the mean memory signal has reached its peak and will therefore fall. Only after this point do the strength and filter distributions begin to re-symmetrise between weak and strong synapses.

3.2 Massed Versus Spaced Repetition With Decay

Although massed and spaced protocols in the absence of a filter decay mechanism could in principle be distinguished by synapses by postulating for example some form of temporal filtering of the perceptron's firing rate, the underlying cause of the perceptron's strong response to a massed protocol compared to a spaced protocol is quite clear. A massed protocol shifts filter states upwards or downwards in a regimented, regular, nearly deterministic manner, allowing little possibility for some degree of re-symmetrisation of filter states between repeated stimuli. A much simpler alternative to a perhaps complicated temporal filter on the perceptron's firing rate is to hypothesise instead that a synapse in some sense becomes dynamically refractory to further filter threshold processes following a strong stimulus. This can be implemented by dynamically modifying the decay rate of filter states when strong stimuli, which activate both neurotransmitter and neuromodulatory systems, occur. In particular, if the filter decay rate increases rapidly following a strong stimulus, then this makes the probability that the filter will reach threshold over some subsequent period of time much lower. This provides the intuition for our proposal that

the filter decay rate $\eta(t)$ is under dynamic control: after a strong stimulus, $\eta(t)$ increases rapidly and remains elevated above basal levels for a period of time, greatly reducing the chance of filter threshold processes during this period. While a massed protocol will be disrupted by this transient increase in $\eta(t)$, a sufficiently spaced protocol will not be, because the stimuli will miss these periods of transiently elevated $\eta(t)$.

With this intuition in mind, we now turn on dynamic decay. For our standard parameters of $\Theta = 8$ and $\rho = 6$, we establish a reference point by searching in parameter space for the values of τ' and $\Delta\eta'$ that maximise the difference between the maximum value of $\mu(t')$ under an at-peak protocol and the corresponding value under a massed protocol. Massed stimulus times are again $t'_i = i$, while for each candidate choice of τ' and $\Delta\eta'$, the at-peak times are again determined numerically and iteratively, with a set $\{t'_0, \ldots, t'_i\}$ inducing the next peak and hence repetition time t'_{i+1}. Having found this reference pair $(\tau', \Delta\eta')$, we also vary them independently three-fold, so that we consider a set of nine pairs $(3^i\tau', 3^j\Delta\eta')$, with $i, j \in \{-1, 0, +1\}$, allowing a comparison between this best at-peak reference point and nearby points in parameter space.

We find this reference point to be given by parameters $\tau' \approx 3.31$ and $\Delta\eta' \approx 0.19$, where we search to two decimal places. Results for the evolution of $\mu(t')$ and $\eta'(t')$ under massed and at-peak protocols for the nine points in parameter space are shown in Fig. 4. Each panel shows the evolution of these two quantities, split into two graphs, for each pair of parameters, as indicated, with the central panel, E, showing the reference pair and the others the three-fold variations around them. For the reference pair, $\mu(t')$ for the massed protocol is considerably suppressed compared to that for the at-peak protocol. The decay rate $\eta'(t')$ accumulates over each massed repetition, with

each successive impulse overcoming any intervening reduction in $\eta'(t')$. In contrast, the spacing between at-peak repetitions is such that $\eta'(t')$ returns close to zero before the next strong stimulus arrives. A non-zero value of $\eta'(t')$ will inevitably affect a spaced protocol, but not as much as a massed protocol. In panels A, D and G, $\Delta\eta'$ is three times smaller than its reference value, but τ' varies, increasing from A to D to G. Although there is some reduction in $\mu(t')$ for the at-peak protocol as τ' increases, the most significant impact is on the massed protocol, with $\mu(t')$ being reduced around three-fold between panels A and G. The same trend is seen for any fixed choice of $\Delta\eta'$ with τ' increasing, as in panels B, E and H or panels C, F and I. When τ' is increased, the window over which $\eta'(t')$ remains transiently elevated increases in length, so that massed repetitions will at some point start to fall within these windows. Of course, if τ' is increased enough, spaced repetitions will also be caught in this manner, so there is an optimal choice of τ' that significantly affects massed protocols without affecting spaced protocols too much. Increasing $\Delta\eta'$ while keeping τ' constant, so moving horizontally between the panels in Fig. 4, again reduces $\mu(t')$ for either protocol. The impact is significant in both cases, but even more so for massed protocols, in which again massed repetition is more likely to fall within the transiently elevated decay window. By the time both $\Delta\eta'$ and τ' have increased three-fold compared to their reference values, $\mu(t')$ for both protocols is essentially completely suppressed, with hardly any response. The pair of values corresponding to the reference $(\Delta\eta', \tau')$ is therefore a balance between suppressing the response to a massed protocol while not suppressing the response to a spaced protocol too much. If both $\Delta\eta'$ and τ' are too small, then neither protocol is affected; while if they are too large, then both are affected. Away from these two extremes in panels A and I, massed protocols are affected more than spaced protocols, in relative terms, with the reference

pair of values used in panel E giving the greatest suppression of the massed compared to the spaced protocol, in absolute terms. For this reference pair of values, a spaced protocol is easily distinguished from a mass protocol simply by placing a threshold on the perceptron's firing rate, because there are very clear margins between these two protocols.

For the choice of the reference parameters used in Fig. 4, we examine the distribution of filter states and the various threshold and splitting probabilities for massed and at-peak protocols in Fig. 5. Fig. 5 should be compared to Fig. 3, for which dynamic filter decay was absent; the formats of these two figures are identical. In the presence of dynamic decay for the massed protocol, we see that while the distribution of filter states for both weak and strong synapses is inevitably skewed in favour of higher rather than lower filter states, the probability of being in the very highest filter states is very much reduced, and indeed becomes increasingly reduced with more repetitions. This of course is a direct consequence of the dynamic decay. While the massed repetition of stimuli shifts the filter distributions upwards, the dynamic decay, whose influence is stronger further away from the zero state, pulls the distributions back to the zero state. For a massed protocol with dynamic decay, the filter distributions for weak and strong synapses are rather similar in overall shape, while without a decay process, in Fig. 3, their shapes are very different. For the at-peak protocol, in contrast, the filter distributions are very similar in overall shape in Figs. 3 and 5, for both weak and strong synapses, indicating that the decay dynamics have much less impact on the at-peak protocol compared to the massed protocol. Further, in Fig. 5B, we explicitly see that for the massed protocol, $\mathsf{Prob}[I = +(\Theta - 1) \,|\, S_i = -1]$ falls with each repetition in the presence of filter decay, while without decay, in Fig. 3B, it increases. Moreover, while the splitting probabilities $\langle \pi_J^+ \rangle_J$ for weak and strong synapses in Fig. 3C are

clearly quite different, in Fig. 5C with decay, they are very similar. Overall, then, a comparison of Figs. 3 and 5 confirms that dynamic filter decay operates precisely as our intuition suggested, by making synapses transiently refractory to filter threshold processes, and that this refractoriness impacts on massed protocols more than on spaced protocols, at least in (large) regions of parameter space.

3.3 Optimally Spaced Repetition Protocols

We have used the at-peak protocol as an example of spaced repetition because the mean memory signal peaks are easy and quick to find, numerically-speaking, and furthermore, the signal peaks present themselves as natural locations at which to seek to reinforce the storage of a memory by repeating strong stimuli. However, the at-peak protocol is not necessarily the optimally spaced protocol, in terms of maximising the difference between the maxima of $\mu(t')$ for spaced and massed stimuli. For a fixed choice of filter size Θ and repetitions ρ, we must search through parameter space to find this optimally spaced protocol. To perform this search, for a fixed pair of parameters $(\tau', \Delta\eta')$, we must search for that set of spaced repetition times $\{t'_1, \ldots, t'_\rho\}$ that maximises this difference. The search space in repetition times is of course ρ-dimensional, but fortunately the search is essentially a convex optimisation problem, so we can use gradient ascent methods to find the optimal times. Each potential step involves numerically re-solving the system in Eq. (16) multiple times, so we restrict to steps of unit size in each time direction for increased speed and tractability. Each such search returns for any given pair $(\tau', \Delta\eta')$ the spaced repetition times that maximise the difference between spaced and massed protocols. For fixed Θ and ρ we must then find the optimal choice of the pair $(\tau', \Delta\eta')$ that maximises this maximum difference. This search is computa-

tionally very intensive and is performed on a cluster supercomputer.

Fig. 6 shows the locations of these optimal points in both the η_0–τ' plane and the $\Delta\eta'$–τ' plane for various values of Θ and ρ. We show the values corresponding to optimally spaced protocols and, for comparison, the values corresponding to at-peak protocols (for which we still search for the best choices of τ' and $\Delta\eta'$, but the at-peak stimulus times are taken as earlier to correspond to each successive peak in $\mu(t')$). We first observe that although there are differences in parameter values between optimal and at-peak stimuli, their values are nevertheless extremely similar. Indeed, the similarity between their values allows us to use the at-peak values as the initial seeds in the search for the optimal values. Their similarity suggests that spaced protocols are not excessively sensitive to the precise choice of the parameters $\Delta\eta'$ and τ', nor to the repetition times $\{t'_1, \ldots, t'_\rho\}$, as we explore later.

As Θ increases, the best choice of τ' increases. This dependence is to be expected, since as Θ increases, the time-to-peak increases, quadratically in Θ from Eq. (8). Thus, the duration of the decay window may be extended as Θ increases, so as to capture more effectively massed protocols without interfering too much with spaced protocols. Conversely, as Θ increases, the best choice of $\Delta\eta'$ or η_0 decreases. Again, this follows from the behaviour of τ'. By reducing $\Delta\eta'$, the impact on spaced stimuli is reduced, while the increasing duration of the decay window ensures that smaller impulses $\Delta\eta'$ accumulate without significant decay during mass repetitions, so that their accumulated impact serves to suppress the response to massed stimuli. As the number of repetitions increases, we see a form of convergence of best parameter values. In general, except at very small ρ for higher Θ, the best choice of $\Delta\eta'$ reduces as ρ increases, again because with more repetitions, the impulses can accumulate. The trend in η_0 does not exhibit the initial non-monotonicity for

smaller ρ and larger Θ. As a compound parameter, η_0 also contains a factor of τ', so it is affected by the change in τ'. As ρ increases, the best choice of τ' decreases, because again the effect of the decay process will accumulate with higher repetitions, even if $\eta(t)$ falls somewhat more between repetitions. In this case, the impact of decay on a spaced protocol will be reduced while continuing to affect a massed protocol. The best choices of parameters are clearer in the η_0–τ' plane, but when shown in the $\Delta\eta'$–τ' log-log plane, we observe that for fixed numbers of repetitions, the best values fall on a very nearly straight line for different choices of Θ, thereby demonstrating power-law behaviour.

Corresponding to the optimal parameters τ' and $\Delta\eta'$ in Fig. 6, we show in graphical form the repetition times $\{t'_1, \ldots, t'_\rho\}$ in Fig. 7 that maximise the difference between spaced and massed protocols. Again we show results for optimal and at-peak protocols. In addition to the basic, unnormalised times t'_i we also show a normalised form, t'_i/Θ^2, where the expected quadratic dependence of the at-peak repetition times on Θ has been removed. For at-peak protocols with fixed Θ, if τ' and $\Delta\eta'$ were held fixed as ρ changes, then the repetition times would not depend on ρ (e.g., t'_1 would not depend on $\rho \geq 1$, etc.) as they are defined as the successive peak times of $\mu(t')$, which would be determined by the fixed choice of τ' and $\Delta\eta'$. However, when we search for the maximising parameters τ' and $\Delta\eta'$ for each choice of ρ, the induced at-peak repetition times depend on ρ because the peak times depend on τ' and $\Delta\eta'$. Therefore, unnormalised and normalised times for optimal *and* at-peak protocols show a systematic dependence on both Θ and ρ. For fixed Θ, more repetitions lower the required repetition times compared to fewer repetitions. This trend is clearer for optimal stimuli than for at-peak stimuli. For at-peak repetition, the normalised repetition times are relatively but not completely insensitive to both Θ and ρ. For example, regardless of Θ, for

at-peak repetition, $t'_3/\Theta^2 \approx 1$ for any choice of $\rho \geq 3$ in this figure, while for optimal repetition, t'_3/Θ^2 shows greater variation with ρ. For $\Theta > 4$, compared to at-peak repetition times the optimal repetition times are earlier, so that for optimal protocols, repetitions occur before the mean memory signal has reached its next peak. The reverse is true for $\Theta \leq 4$ (data not shown), with optimal repetitions occurring later than at-peak repetitions, but small threshold filters are not typical.

We further illustrate these trends in Fig. 8, in which we plot $\mu(t')$ for massed, optimal and at-peak protocols for different choices of Θ and ρ, as indicated. Each panel uses the optimal parameters $(\tau', \Delta\eta')$ for the given choice of Θ and ρ, and the at-peak and massed repetition times are generated in the usual manner. In each panel, we normalise $\mu(t')$ relative to the maximum value of the optimal signal to facilitate easy comparison within and between panels. By definition of the optimal protocol, its $\mu(t')$ always attains the normalised maximum of unity. We clearly see that the optimal maximum always occurs before the at-peak maximum, with the optimal repetitions always occurring before the corresponding at-peak repetitions (at least here, for $\Theta > 4$). For the smaller choices of Θ or ρ shown, the optimal and at-peak repetition and maximum times are similar. Only for larger choices of Θ and ρ do we begin to see a clear separation of repetition times for these two protocols. What is striking, however, is that despite the optimisation of the difference between spaced and massed maxima, optimal and at-peak protocols induce rather similar, sometimes very similar, maximum values of $\mu(t')$. Even for $\Theta = 16$ and $\rho = 12$ in panel I, the at-peak maximum is around 92% of the optimal maximum. These results confirm that although optimal and at-peak repetition times can be quite different, the particular choice of a spaced stimulus does not require exquisite tuning, since both optimal and at-peak repetition are clearly and robustly dis-

tinguishable from massed repetition, in terms of attaining similar maximum values of $\mu(t')$. We will see next that not only do the spaced repetition times $\{t'_1, \ldots, t'_\rho\}$ not need exquisite tuning, but neither do the parameter τ' and $\Delta\eta'$ need tuning to their optimal values.

Optimal choices of parameters have been determined here on the basis of comparing massed and spaced protocols with identical numbers of repetitions. This leaves open the question of whether massed protocols with more repetitions could increase the massed maximum value of $\mu(t')$ so that it becomes comparable to or even exceeds that for spaced protocols with fewer repetitions. That is, can we simply repeat a massed stimulus more and more times until its peak exceeds a spaced stimulus with a fixed number of repetitions? From our earlier considerations, it is clear that more massed repetitions suppress the probability that a filter may subsequently reach threshold (e.g. Fig. 5B). We therefore cannot simply repeat a massed stimuli more times and drive its peak higher. The filter decay process ensures that this is not a viable option. The only way to increase the maximum of $\mu(t')$ is to space the stimuli apart, so that they miss the decay windows. The use of equal numbers of repetitions to compare spaced and massed protocols and obtain optimal parameters is therefore valid.

3.4 Suboptimal but Good Enough Spaced Repetition

The results in Fig. 6 give the optimal values of τ' and $\Delta\eta'$ (or η_0) that maximise the difference between spaced and massed protocols. While a single synapse may be expected to be able to set τ' and $\Delta\eta'$ in a manner that depends on its filter threshold size, it of course cannot and does not have access to the number of times that a future strong stimulus will be presented. Setting τ' and $\Delta\eta'$ as a function of the number of repetitions, ρ, to achieve optimal behaviour is

therefore not an available option. But can suboptimal behaviour that is good enough, in the sense of nevertheless robustly distinguishing between spaced and massed protocols, be achieved? That is, can the parameters τ' and $\Delta\eta'$ be set in a Θ-dependent but ρ-independent manner and still achieve robust discrimination?

To answer this question, first in Fig. 9 we plot in the η_0–τ' plane 90% contours around the optimal points shown in Fig. 6. These contours correspond to values of τ' and η_0 for which the difference between spaced and massed protocols is 90% of the maximum possible difference for optimal choices of τ' and η_0. (We note that for each point in the η_0–τ' plane, we always search for the repetition times $\{t'_1, \ldots, t'_\rho\}$ that maximise this difference, so these times depend explicitly on τ' and η_0; they have not been fixed at the repetition times corresponding the optimal choices of τ' and η_0.) For comparison we again show results for both optimal and at-peak protocols, in panels A and B, respectively. We plot contours for various choices of Θ as shown, excluding some values only to avoid excessive clutter, and for all values of ρ that we have examined, from $\rho = 2$ to $\rho = 12$. For each of the smaller values of Θ shown, the set of contours for all values of ρ always enclose a common, core region. Inside this region, the difference between spaced and massed protocols is always in excess of 90% of optimal, so that regardless of the actual choice of ρ, we are assured of at least 90% optimal difference in this region. For larger values of Θ, the contours initially shift rapidly for small ρ, but if we exclude the $\rho = 2$ case, there is still always a sizeable, common core with at least 90% optimal difference; including $\rho = 2$ restricts us to a relatively small region. In these upper panels, we also indicate the average values of the optimal choices of η_0 and τ' for a given value of Θ, with the average taken over ρ. We consider either all examined values of ρ or all values excluding $\rho = 2$ and $\rho = 3$, which values tend to distort the

general trends. These average values, plotted explicitly as a function of Θ in panels C and D, mostly sit at the centres of the corresponding core regions in panels A and B. They are therefore candidate choices of τ' and η_0 for a synapse to achieve at least 90% optimal discrimination between spaced and massed repetition, regardless of the number of repetitions. For τ', we see the power-law dependence on Θ in panels C and D. For η_0, the power-law behaviour is less clear particularly for smaller Θ, as discussed above, but had we instead plotted $\Delta\eta'$, its power-law behaviour would have been clearer. Indeed, we have used η_0 rather than $\Delta\eta'$ in this figure only for greater clarity (cf. Fig. 6). Had we plotted only $\Delta\eta'$, the inclusion or exclusion of small ρ in the averaging would have been entirely moot, as the two corresponding average values of $\Delta\eta'$ would have been much more similar for larger Θ than those for η_0 with its extra factor of τ'. Fitting the average values (over all ρ) of τ' and $\Delta\eta'$ to power laws in Θ for $\Theta = 5, \ldots, 16$ (so excluding the noisy, unrepresentative $\Theta = 4$ case: see Fig. 6), for optimal spacing we obtain $\tau'(\Theta) \approx 0.063 \times \Theta^{1.87}$ (so slightly sub-quadratic in Θ) and $\Delta\eta'(\Theta) \approx 52.1/\Theta^{2.65}$, and for at-peak spacing we obtain $\tau'(\Theta) \approx 0.052 \times \Theta^{1.99}$ (almost perfectly quadratic in Θ) and $\Delta\eta'(\Theta) \approx 53.9/\Theta^{2.67}$. Finally, we remark that the various results for optimal and at-peak repetition in this figure are very similar, as the fits also confirm, again reflecting that the discrimination between spaced and massed protocols is relatively insensitive to the precise details of the spaced repetition times, whether optimal or at-peak.

In order that this ability to select regions of parameter space with at least 90% performance compared to optimal, regardless of ρ and without excessive sensitivity to the repetition times, is actually useful, we must also exhibit the dynamics of $\mu(t')$ with parameters taken from these 90% contours and show that the reduction is not significant. We do this in Fig. 10, in which we take

$\Theta = 8$ and $\rho = 6$ as our standard values, take the corresponding optimal values of τ' and η_0, and then consider eight points on the 90% contour around this optimal point. The dynamics of $\mu(t')$ and $\eta'(t')$ for these nine points in τ'–η_0 parameter space, including the optimal point, are shown in the nine panels of Fig. 10, with the central panel E corresponding to the optimal point. Inevitably, as the parameters vary around the 90% contour, $\mu(t')$ varies for both massed and spaced protocols. However, for the optimal point, we have maximised the maximum difference between the spaced and massed signals. The absolute changes in $\mu(t')$ are therefore not important. Rather, what is important is that the differences between the spaced and massed maxima of $\mu(t')$ do not vary much around the 90% contour relative to the optimal point. We have therefore also plotted a rectified, relative form of the spaced mean memory signal $\mu(t')$ in which we subtract from it the maximum value of the massed mean memory signal, truncating negative values where necessary. We see that the maximum of this rectified, relative signal is in all panels around 0.2, in fact varying between 0.194 and 0.215 while the maximum of the spaced signal varies between 0.24 and 0.35 while that for the massed signal varies between 0.04 and 0.15. Thus, the suboptimal parameter choices around this 90% contour discriminate between massed and spaced protocols almost as well as the optimal parameter choice. Hence, we suffer very little deterioration in performance by relaxing a strict optimality requirement. A synapse may therefore very effectively discriminate between spaced and massed protocols, regardless of the number of repetitions, to which it cannot have prior access, and also without extreme sensitivity to the precise timing of the spaced repetitions, also to which it cannot have prior access.

3.5 Typical Versus Particular Behaviour

We have so far examined the typical or average dynamics of $h(t)$, via $\mu(t)$, where we average over all possible realisations of the tracked and repeated memory $\boldsymbol{\xi}^0$ and the non-tracked and unrepeated memories $\boldsymbol{\xi}^\alpha$, $\alpha \geq 1$. This allows us to examine the difference between massed and spaced protocols without having to consider noise or fluctuations. However, real memory systems store real memories, which of course are always particular realisations. We must therefore show that the mean behaviour is completely representative, in the sense that particular realisations do not vary from it so much as to make the above analysis of the mean behaviour redundant. To this end, we run simulations of our model according to methods described at length elsewhere (Elliott & Lagogiannis, 2012; Elliott, 2016a). We may average over multiple simulations to obtain average statistics to validate our results for $\mu(t)$, but we may also use single simulations to obtain results for particular realisations of $h(t)$. We do this for $\Theta = 8$ and $\rho = 6$, as standard, taking the optimal parameters $\tau' = 3.16$ and $\eta_0 = 0.59$ (see Fig. 10E), and using the optimal repetition times of $t'_i \in \{21, 37, 51, 64, 77, 90\}$, $i = 1, \ldots, 6$, with of course $t'_0 = 0$ as the initial storage time for $\boldsymbol{\xi}^0$. For each simulation, all the memories have components ξ_i^α drawn randomly from $\{-1, +1\}$ and the times for the storage of $\boldsymbol{\xi}^\alpha$ for $\alpha \geq 1$ are also drawn according to their Poisson distribution. For examining $\mu(t)$ we have not so far had to specify N, the number of synapses, as $\mu(t)$ is an average over N identically distributed random variables. However, for running simulations, we must specify N, which also sets the scale for the fluctuations in $h(t)$, measured here by its variance, $\sigma^2(t)$. Results for $N = 10^3$, $N = 10^4$ and $N = 10^5$, spanning a biologically plausible range of values, are shown in Fig. 11. We plot the mean $\mu(t')$, here determined from simulations

but up to simulation noise identical to the numerical results discussed above; the one and two standard deviation regions around it (1σ and 2σ); and for each choice of N, four completely representative examples of particular realisations of $h(t')$. We see that the particular realisations of $h(t')$ closely follow the overall trend in $\mu(t')$, so that the realisations robustly increase with each repetition of the tracked memory and reach a maximum similar to that for the mean. The fluctuations in the realisations are mostly confined to the 1σ and 2σ regions around the mean, with these regions being quite narrow. Up to issues of correlations between the strengths of pairs of synapses (Elliott & Lagogiannis, 2012), which we can largely ignore here, the variance in synaptic strength is around $1/N$, and even for $N = 10^3$, the mean at its maximum swamps the standard deviation. Hence, the memory signal for any particular memory behaves similarly to the mean signal and is subject to fluctuations that are sufficiently small that if we can robustly discriminate between massed and spaced stimuli at the level of the mean, then we can also discriminate between them in real scenarios with real, particular rather than typical memories. A simple threshold on the perceptron's firing rate will work for both average dynamics and for particular realisations.

4 Discussion

Integrate-and-express models of synaptic plasticity propose that synapses act as low-pass filters, integrating plasticity induction signals before expressing synaptic plasticity. Such models powerfully suppress fluctuations in synaptic strength and possess very rich dynamics, including an initially rising memory signal driven by the ongoing storage of other memories, which in non-integrative models reduces rather than enhances the memory signal. This rise

in the memory signal to a maximum suggests a natural time at which to repeat the memory for re-storage, hinting at the spacing effect. To fully explain the spacing effect with synaptic integration, we have postulated that filter decay dynamics are regulated by coincident neurotransmitter and neuromodulatory activity. With such dynamic regulation, a synaptic filter provides a natural and robust mechanism for a synapse to discriminate between massed and spaced repetition of strong stimuli, suppressing the response to the former while maintaining the response to the latter. The two key parameters that characterise the regulated filter decay process can be interpreted as the duration and amplitude of activity in signalling cascades downstream from cAMP production resulting from neuromodulation, perhaps specifically the MAPK pathway and the phosphatase SHP2 that is known to modulate the spacing effect (Pagani *et al.*, 2009).

We have focused exclusively on the spacing effect in the transition from early- to late-phase plasticity, where we have understood the spacing effect to be the capacity of a single synapse to discriminate between massed and spaced repetition protocols. Evidence does indeed support the view that the spacing effect is a property of single synapses, giving specificity and precision to structural changes (Martin *et al.*, 2017). In particular, we have implicitly assumed that the capacity to discriminate between massed and spaced repetition must be both logically and biologically prior to the capacity to convert early- into late-phase plasticity: the former is a necessary (but not sufficient) condition for the latter. Thus, we have taken the latter as dissociable from the former, so that we can set to one side, for future study, the mechanistic underpinnings of the latter. Therefore, it has not been necessary for us to consider issues of synaptic tagging, synaptic capture and cross-capture, and late associativity (Frey & Morris, 1997; Frey & Morris, 1998; Sajikumar & Frey, 2004; Rey-

mann & Frey, 2007; Frey & Frey, 2008; Wang & Morris, 2009). Many of these processes are heterosynaptic, extending into dendritic compartments, and ultimately involving the nucleus via, for example, the CREB pathway. To consider the process by which single synapses convert early- into late-phase plasticity, after determining that a stimulus is suitably spaced rather than massed, would necessarily entail moving beyond the purely single synapse level. A few models have considered synaptic tagging (Clopath *et al.*, 2008; Barrett *et al.*, 2009; Päpper *et al.*, 2011; Ziegler *et al.*, 2015), but they are largely phenomenological in character, fitting model parameters to experimental data in state-based models without providing much explanatory or mechanistic insight.

Many models have considered memory formation with discrete-strength synapses in a feedforward or recurrent network setting (Tsodyks, 1990; Amit & Fusi, 1994; Fusi *et al.*, 2005; Leibold & Kempter, 2006; Barrett & van Rossum, 2008; Huang & Amit, 2010; Elliott & Lagogiannis, 2012; Lahiri & Ganguli, 2013). These models overcome the catastrophic forgetting of the classic Hopfield network (Hopfield, 1982) by allowing the forgetting of older memories as newer ones are stored, turning them into so-called "palimpsest" memories (Nadal *et al.*, 1986; Parisi, 1986). Common to all these palimpsest models, including that studied here, is that memories are essentially stored as transients on an equilibrium distribution, and the transients inevitably die away as the system returns to equilibrium. In all non-integrative models (all those but our own), this process of returning to equilibrium occurs immediately after the storage of a memory, but in integrative models, it starts only after the memory signal has reached its peak, with the system moving further away from equilibrium prior to peak signal. Memory lifetimes are determined by the longevity of these transients, and may be defined in a variety of ways, including via a signal-to-noise ratio (Tsodyks, 1990), the completely equivalent

"ideal observer" approach (Fusi *et al.*, 2005; Lahiri & Ganguli, 2013; Elliott, 2016b), and a mean first passage time method (Elliott, 2014). Recently, using this last method, we showed that in fact the storage capacity of a palimpsest memory system is fatally compromised when a perceptron's firing threshold is not *precisely* tuned to match its equilibrium mean memory signal (Elliott, 2014; Elliott, 2017). Further, even when this unrealistic tuning is assumed, many models require that a neuron can discriminate changes in its membrane potential far below the roughly 0.5mV level of noise level inherent in membrane potentials (Sigworth, 1980). Such noise completely destroys these models' performance.

It is for these reasons that we have not examined memory lifetimes, via any definition, in the current study. Regardless of spaced repetition of strong stimuli, memories will inevitably die away, sooner or later, in palimpsest models, whether integrative or non-integrative. The difference between short- and long-term memory is not one of degree, but one of kind. Some new, fundamental idea must therefore be introduced into these models so that they become viable rather than merely toy models of biological memory systems. That idea, of course, must be precisely the transition from early- to late-phase plasticity, essentially locking synapses into non-labile states, preventing rather than merely delaying a return to equilibrium. The signal to lock may simply be when a neuron's firing rate exceeds some threshold, so that a back-propagating action potential robustly invades the dendritic tree, perhaps initiating the process of setting synaptic tags at those synapses that participated in firing the neuron strongly. We have begun this reformulation here, by showing that a single synapse can robustly discriminate between massed and spaced stimuli. In future work we must consider the consequences of locking synapses into non-labile states. For example, such a memory system will immediately break

the intimate connection between memory capacity and memory longevity in palimpsest memory. The memory system will slowly reach maximum capacity, although presumably in a non-catastrophic manner, of simply being unable to store further memories as it reaches saturation. Short-term and long-term memory systems are different, with hippocampal-neocortical interactions being vital in transferring memories to long-term systems (Wang & Morris, 2009), so saturation is not necessarily a fundamental problem. But it is necessary to understand how such a memory system could continue to function, perhaps by unlocking synapses involved in the storage of memories that have since been transferred to other networks, or perhaps by ongoing synaptogenesis, making new synapses available, or even, in some structures, by ongoing neurogenesis, making entirely new sub-assemblies available. It will be fascinating to explore these rich themes in extensions of the present work.

5 Conclusion

We have shown that when synaptic filtering is dynamically coupled to neuromodulatory signalling, a natural and powerful explanation of the spacing effect in the transition from short- to long-term memory emerges. By thoroughly exploring the model's parameter space, we have found large regions where single synapses can robustly distinguish between massed and spaced protocols, and other regions where such discrimination breaks down. Although synaptic filtering was originally postulated as a mechanism for the control of potentially destabilising fluctuations in synaptic strength, its unbidden and automatic capacity also to provide a novel and deep explanation of the spacing effect lends support to our proposal that single synapses should be regarded as plasticity induction signal processors.

Acknowledgements: I acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work.

References

Alon, U. 2006. *An Introduction to Systems Biology: Design Principles of Biological Circuits*. London: Chapman & Hall, CRC.

Amit, D.J., & Fusi, S. 1994. Learning in neural networks with material synapses. *Neural Comput.*, **6**, 957–982.

Andersen, P., Morris, R.G.M., Amaral, D., Bliss, T.V.P., & O'Keefe, J. (eds). 2007. *The Hippocampus Book*. Oxford: Oxford University Press.

Asrican, B., Lisman, J., & Otmakhov, N. 27. Synaptic strength of individual spines correlates with bound Ca^{2+}-calmodulin-dependent kinase II. *J. Neurosci.*, **2007**, 14007–14011.

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. *Proc. Natl. Acad. Sci. U.S.A.*, **102**, 14434–14439.

Bailey, C.H., Montarolo, P., Chen, M., Kandel, E.R., & Schacher, S. 1992. Inhibitors of protein and RNA synthesis block structural changes that accompany long-term heterosynaptic plasticity in Aplysia. *Neuron*, **9**, 749–758.

Barco, A., Alarcon, J.M., & Kandel, E.R. 2002. Expression of constitutively active CREB protein facilitates the late phase of long-term potentiation by enhancing synaptic capture. *Cell*, **108**, 689–703.

Barrett, A.B., & van Rossum, M.C.W. 2008. Optimal learning rules for discrete synapses. *PLoS Comput. Biol.*, **4**, e1000230.

Barrett, A.B., Billings, G.O., Morris, R.G.M., & van Rossum, M.C.W. 2009. State based model of long-term potentiation and synaptic tagging and capture. *PLoS Comput. Biol.*, **5**, e1000259.

Bartol, T.M., Bromer, C., Kinney, J., Chirillo, M.A., Bourne, J.N., Harris, K.M., & Sejnowski, T.J. 2015. Nanoconnectomic upper bound on the variability of synaptic plasticity. *eLife*, **4**, e10778.

Bartsch, D., Ghirardi, M., Skehel, P.A., Karl, K.A., Herder, S.P., Chen, M., Bailey, C.H., & Kandel, E.R. 1995. Aplysia CREB2 represses long-term facilitation: Relief of repression converts transient facilitation into long-term functional and structural change. *Cell*, **83**, 979–992.

Beck, C.D., Schroeder, B., & Davis, R.L. 2000. Learning performance of normal and mutant *Drosophila* after repeated conditioning trials with discrete stimuli. *J. Neurosci.*, **20**, 2944–2953.

Bhalla, U.S., & Iyengar, R. 1999. Emergent properties of networks of biological signaling pathways. *Science*, **283**, 381–387.

Bliss, T.V.P., & Lømo, T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the unanaesthetized rabbit following stimulation of the perforant path. *J. Physiol.*, **232**, 331–356.

Carew, T.J., Pinsker, H.M., & Kandel, E.R. 1972. Long-term habituation of a defensive withdrawal reflex in Aplysia. *Science*, **175**, 451–454.

Clopath, C., Ziegler, L., Vasilaki, E., Büsing, L., & Gerstner, W. 2008. Tag-trigger-consolidation: A model of early and late long-term-potentiation and depression. *PLoS Comput. Biol.*, **4**, e1000258.

Davis, R.L. 2005. Olfactory memory formation in *Drosophila*: From molecular to systems neuroscience. *Annu. Rev. Neurosci.*, **28**, 275–302.

Ebbinghaus, H. 1885. *Memory: A Contribution to Experimental Psychology.* New York, NY: Dover Publications.

Eichenbaum, H., & Cohen, N.J. 2001. *From Conditioning to Conscious Recollection.* Oxford: Oxford University Press.

Elliott, T. 2008. Temporal dynamics of rate-based plasticity rules in a stochastic model of spike-timing-dependent plasticity. *Neural Comput.*, **20**, 2253–2307.

Elliott, T. 2014. Memory nearly on a spring: A mean first passage time approach to memory lifetimes. *Neural Comput.*, **26**, 1873–1923.

Elliott, T. 2016a. The enhanced rise and delayed fall of memory in a model of synaptic integration: Extension to discrete state synapses. *Neural Comput.*, **28**, 1927–1984.

Elliott, T. 2016b. Variations on the theme of synaptic filtering: A comparison of integrate-and-express models of synaptic plasticity for memory lifetimes. *Neural Comput.*, **28**, 2393–2460.

Elliott, T. 2017. Mean first passage memory lifetimes by reducing complex synapses to simple synapses. *Neural Comput.*, **29**, 14681527.

Elliott, T., & Lagogiannis, K. 2009. Taming fluctuations in a stochastic model of spike-timing-dependent plasticity. *Neural Comput.*, **21**, 3363–3407.

Elliott, T., & Lagogiannis, K. 2012. The rise and fall of memory in a model of synaptic integration. *Neural Comput.*, **24**, 2604–2654.

Ferrell, J.E. 1996. Tripping the switch fantastic: how a protein kinase cascade can convert graded inputs into switch-like outputs. *Trends Biochem. Sci.*, **21**, 460–466.

Frey, S., & Frey, J.U. 2008. 'Synaptic tagging' and 'cross-tagging' and related associative reinforcement processes of functional plasticity as the cellular basis for memory formation. *Prog. Brain Res.*, **169**, 117–143.

Frey, U., & Morris, R.G.M. 1997. Synaptic tagging and long-term potentiation. *Nature*, **385**, 533–536.

Frey, U., & Morris, R.G.M. 1998. Synaptic tagging: Implications for the late maintenance of hippocampal long-term potentiation. *Trends Neurosci.*, **21**, 181–188.

Frey, U., Huang, Y.-Y., & Kandel, E.R. 1993. Effects of cAMP simulate a late stage of LTP in hippocampal CA1 neurons. *Science*, **260**, 1661–1664.

Fusi, S., Drew, P.J., & Abbott, L.F. 2005. Cascade models of synaptically stored memories. *Neuron*, **45**, 599–611.

Harris, K.M., & Stevens, J.K. 1989. Dendritic spines of CA1 pyramidal cells in the rat hippocampus: serial electron microscopy with reference to their biophysical characteristics. *J. Neurosci.*, **9**, 2982–2887.

Hawkins, R.D., Abrams, T.W., Carew, T.J., & Kandel, E.R. 1983. A cellular mechanism of classical conditioning in *Aplysia*: Activity dependent amplification of presynaptic facilitation. *Science*, **219**, 400–405.

Hawkins, R.D., Kandel, E.R., & Bailey, C.H. 2006. Molecular mechanisms of memory storage in *Aplysia*. *Biol. Bull.*, **210**, 174–191.

Heisenberg, M. 2003. Mushroom body memoir: From maps to models. *Nature Rev. Neurosci.*, **4**, 266–275.

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. *Proc. Natl. Acad. Sci. U.S.A.*, **79**, 2554–2558.

Huang, Y., & Amit, Y. 2010. Precise capacity analysis in binary networks with multiple coding level inputs. *Neural Comput.*, **22**, 660–688.

Huang, Y.-Y., & Kandel, E.R. 1995. D1/D5 receptor agonists induce a protein synthesis-dependent late potentiation in the CA1 region of the hippocampus. *Proc. Natl. Acad. Sci. U.S.A.*, **92**, 2446–2450.

Krug, M., Lössner, B., & Ott, T. 1984. Anisomycin blocks the late phase of long-term potentiation in the dentate gyrus of freely moving rats. *Brain Res. Bull.*, **13**, 39–42.

Lahiri, S., & Ganguli, S. 2013. A memory frontier for complex synapses. *Pages 1034–1042 of:* Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., & Weinberger, K.Q. (eds), *Advances in Neural Information Processing Systems 26*. Cambridge, MA: MIT Press.

Leibold, C., & Kempter, R. 2006. Memory capacity for sequences in a recurrent network with biological constraints. *Neural Comput.*, **18**, 904–941.

Lisman, J., & Zhabotinsky, A.M. 2001. A model of synaptic memory: A CaMKII/PP1 switch that potentiates transmission by organizing an AMPA receptor anchoring assembly. *Neuron*, **31**, 191–201.

Lynch, G.S., Dunwiddie, T., & Gribkoff, V. 1977. Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature*, **266**, 737–739.

Malenka, R.C., Kauer, J.A., Perkel, D.J., Mauk, M.D., Kelly, P.T., Nicoll, R.A., & Waxham, M.N. 1989. An essential role for postsynaptic calmodulin and protein kinase activity in long-term potentiation. *Nature*, **340**, 554–557.

Manahan-Vaughan, D., Kulla, A., & Frey, J.U. 2000. Requirement of translation but not transcription for the maintenance of long-term depression in the CA1 region of freely moving rats. *J. Neurosci.*, **20**, 8572–8576.

Martin, A.S., Rela, L., Gelb, B., & Pagani, M.R. 2017. The spacing effect for structural synaptic plasticity provides specificity and precision in plastic changes. *J. Neurosci.*, **37**, 4992–5007.

Miller, P., Zhabotinsky, A.M., Lisman, J.E., & Wang, X.-J. 2005. The stability of a stochastic CaMKII switch: Dependence on the number of enzyme molecules and protein turnover. *PLoS Biol.*, **3**, 0705.

Montgomery, J.M., & Madison, D.V. 2004. Discrete synaptic states define a major mechanism of synapse plasticity. *Trends Neurosci.*, **27**, 744–750.

Mulkey, R.M., Herron, C.E., & Malenka, R.C. 1993. An essential role for protein phosphatases in hippocampal long-term depression. *Science*, **261**, 1104–1107.

Nadal, J.P., Toulouse, G., Changeux, J.P., & Dehaene, S. 1986. Networks of formal neurons and memory palimpsests. *Europhys. Lett.*, **1**, 535–542.

Nusser, Z., Lujan, R., Laube, G., Roberts, J.D., Molnar, E., & Somogyi, P.

1998. Cell type and pathway dependence of synaptic AMPA receptor number and variability in the hippocampus. *Neuron*, **21**, 545–559.

O'Connor, D.H., Wittenberg, G.M., & Wang, S.S.-H. 2005. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proc. Natl. Acad. Sci. U.S.A.*, **102**, 9679–9684.

Pagani, M.R., Oishi, K., Gelb, B.D., & Zhong, Y. 2009. The phosphatase SHP2 regulates the spacing effect for long-term memory induction. *Cell*, **139**, 186–198.

Päpper, M., Kempter, R., & Leibold, C. 2011. Synaptic tagging, evaluation of memories, and the distal reward problem. *Learn. and Mem.*, **18**, 58–70.

Parisi, G. 1986. A memory which forgets. *J. Phys. A: Math. and Gen.*, **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. *Proc. Natl. Acad. Sci. U.S.A.*, **95**, 4732–4737.

Pi, H.J., & Lisman, J.E. 2008. Coupled phosphatase and kinase switches produce the tristability required for long-term potentiation and long-term depression. *J. Neurosci.*, **28**, 13132–13138.

Reymann, K.G., & Frey, J.U. 2007. The late maintenance of hippocampal LTP: Requirements, phases, 'synaptic tagging', 'late-associativity' and implications. *Neuropharm.*, **52**, 24–40.

Roberson, E.D., English, J.D., & Sweatt, J.D. 1996. A biochemist's view of long-term potentiation. *Learn. and Mem.*, **3**, 1–24.

Sajikumar, S., & Frey, J.U. 2004. Late-associativity, synaptic tagging, and the role of dopamine during LTP and LTD. *Neurobiol. Learn. and Mem.*, **82**, 12–25.

Sigworth, F.J. 1980. The variance of sodium current fluctuations at the node of Ranvier. *J. Physiol.*, **307**, 97–129.

Sutton, M.A., Ide, J., Masters, S.E., & Carew, T.J. 2002. Interaction between amount and pattern of training in the induction of intermediate- and long-term memory for sensitization in Aplysia. *Learn. and Mem.*, **9**, 29–40.

Tsodyks, M.V. 1990. Associative memory in neural networks with binary synapses. *Mod. Phys. Lett. B*, **4**, 713–716.

Tully, T., Preat, T., Boynton, S.C., & Del Vecchio, M. 1994. Genetic dissection of consolidated memory in *Drosophila*. *Cell*, **79**, 35–47.

Wang, S.H., & Morris, R.G.M. 2009. Hippocampal-neocortical interactions in memory formation, consolidation, and reconsolidation. *Annu. Rev.Psych.*, **61**, 49–69.

Yin, J.C.P., Del Vecchio, M., Zhou, H., & Tully, T. 1995. CREB as a memory modulator: induced expression of a dCREB2 activator isoform enhances long-term memory in drosophila. *Cell*, **81**, 107–115.

Ziegler, L, Zenke, F., Kastner, D.B., & Gerstner, W. 2015. Synaptic consolidation: From synapses to behavioral modeling. J. Neurosci., **35**, 1319–1334.

Figure Captions

Figure 1: Graphical representation of all allowed transitions between strength and filter states encoded in the transition matrix $r\,\mathbb{G} + \eta(t)\,\mathbb{D}$. Each shaded rectangle represents a definite strength state, $\Omega_1, \ldots, \Omega_\nu$, and each circle enclosing a number represents a definite filter state, $-(\Theta - 1), \ldots, +(\Theta - 1)$. Transitions between states are indicated by directed arrows labelled by their rates, $g_+ r$ and $g_- r$ for processes driven by potentiating and depressing induction signals, and rates proportional to $\eta(t)$ for decay processes (the dependence of η on t has been suppressed for convenience). We stress that although the impression may be gained from this figure that a synapse possesses ν filters, this is not the case. A synapse possesses a single filter, and the indicated transitions are between *joint* strength and filter states.

Figure 2: Comparison of different stimulus protocols in the absence a filter decay process. The mean memory signal $\mu(t')$ as a function of t' is plotted for massed (A), at-peak (C) and randomly-timed (E) stimuli, for $\Theta = 8$. The enveloping solid line in each case shows the cumulative response to 12 successive repetitions (a total of 13 presentations of $\boldsymbol{\xi}^0$) at the times indicated by the arrows along the t'-axis. Also shown, in alternating line styles for clarity, are what the responses would have been had the number of repetitions been only $1, \ldots, 11$, from bottom to top in these three panels. The maximum responses $\max_t \mu(t')$ plotted against the number of repetitions for different choices of Θ are shown for massed (B) and at-peak (D) stimuli, with a direct comparison between massed and at-peak stimuli for particular choices of Θ shown in panel F.

Figure 3: Distributions of filter states and associated threshold and splitting

probabilities in the absence of a filter decay process, for $\Theta = 8$ and $\rho = 6$. (A) For massed and at-peak protocols, we show histograms of the probability that a synapse of lowest strength ($S_i = -1$, "weak") and highest strength ($S_i = +1$, "strong") is in any given filter state. A total of nine time points are taken, corresponding to the seven presentation times for $\boldsymbol{\xi}^0$ ("Rep0", ..., "Rep6"), the time at which $\mu(t')$ reaches peak ("1×Peak"), and twice that time ("2×Peak"). (B) The threshold probabilities $\mathsf{Prob}[J = \pm 7 \,|\, S = \mp 1]$ plotted against time t' for the same massed and at-peak protocols in panel A; we use $1 + t'$ on the abscissa so as not to over-emphasise small times on a logarithmic scale. These probabilities indicate the probability that a weak synapse will become strong and *vice versa*. (C) The corresponding splitting probabilities $\langle \pi_J^+ \rangle_J$ for weak and strong synapses indicate the probability that the synapse will reach the upper filter threshold from any filter state, assuming no further repetitions of $\boldsymbol{\xi}^0$ occur from that time.

Figure 4: Evolution of $\mu(t')$ and $\eta'(t')$, shown separately in upper and lower graphs, for massed (dashed lines) and at-peak (solid lines) protocols for $\Theta = 8$ and $\rho = 6$, in the presence of filter decay for various choices of the parameters τ' and $\Delta\eta'$ as indicated in each panel. In panel E, the reference set of values $(\tau', \Delta\eta') = (3.31, 0.19)$ is taken. This choice maximises the difference between the maximum values of $\mu(t')$ for the spaced and massed protocols. The other panels show results for three-fold variations around these reference parameters.

Figure 5: Distributions of filter states and associated threshold and splitting probabilities, for $\Theta = 8$, $\rho = 6$, $\tau' = 3.31$ and $\Delta\eta' = 0.19$, corresponding to the reference point in Fig. 4E. The format of this figure is otherwise identical to Fig. 3.

Figure 6: Parameters τ' and either η_0 or $\Delta\eta'$ ($\eta_0 = \tau'\Delta\eta'$) that maximise the difference between the maxima of $\mu(t')$ for spaced and massed protocols. Results are shown for optimally spaced protocols in which we search numerically for the repetition times $\{t'_1, \ldots, t'_\rho\}$ and the filter decay parameters that maximise this difference (panels A and C), and for at-peak spaced protocols in which repetitions are always at successive signal peaks but we search for the maximising filter decay parameters (panels B and D). Results are shown in the η_0–τ' plane (panels A and B) or $\Delta\eta'$–τ' plane (panels C and D), for different choices of Θ and ρ, with each point on the "mesh" of points corresponding to the maximising set of decay parameters; Θ varies from 4 to 16 and ρ from 2 to 12, each increasing in the direction of its corresponding arrow.

Figure 7: Graphical representation of the repetition times $\{t'_1, \ldots, t'_\rho\}$ at which the tracked memory should be re-presented for different choices of Θ and ρ, for optimally (top row) and at-peak (bottom row) spaced protocols. For each choice of Θ and ρ, the best choice of parameters τ' and $\Delta\eta'$ is used, so that the optimal times are numerically determined, and the at-peak times are as usual defined by successive peaks. Times are shown in their unnormalised forms t'_i (left-hand column) and normalised with respect to Θ, t'_i/Θ^2 (right-hand column). For each shown choice of Θ, each line denotes a particular value of ρ (increasing from 2 to 12 vertically for each Θ), with the successive repetition times shown as circles along the line.

Figure 8: Comparison of optimal, at-peak and massed repetition protocols for different choices of Θ and ρ, as indicated. In each panel, the parameters τ' and $\Delta\eta'$ are the optimal choices corresponding to the optimal repetition times. The

at-peak repetition times are obtained, using these parameters, in the standard manner, from each successive peak. The open (optimal) and closed (at-peak) circles indicate the times $\{t'_0, \ldots, t'_\rho\}$ for the two spaced protocols. In each panel, we have normalised all forms of $\mu(t')$ by the maximum value of the optimal signal, so that the optimal signal reaches a maximum of unity.

Figure 9: The 90% contours around the maximising parameter points. In panels A (optimal) and B (at-peak) each closed loop is a contour in the η_0–τ' plane corresponding to a maximum difference between spaced and massed protocols of exactly 90% of that for the maximising values of η_0 and τ' (for the given choices of Θ and ρ). Different values of Θ are shown, with each cluster of contours corresponding to the same value of Θ. Alternating line styles are used as Θ increases for clarity; only a few values of Θ are shown to avoid clutter. In each cluster for a fixed value of Θ, each loop shows the contour for a different choice of ρ, with the contours shifting in the η_0–τ' plane in the direction indicated by the arrow as ρ increases from 2 to 12. The circles denote, for fixed Θ, the average values of the optimising values of η_0 and τ', averaged over all choices of ρ; the squares denote these same averages, but excluding the $\rho = 2$ and $\rho = 3$ cases. These average values are explicitly plotted against Θ in panels C (optimal) and D (at-peak).

Figure 10: Suboptimal performance on 90% contours in relation to optimal performance. For $\Theta = 8$ and $\rho = 6$, we consider the optimal parameters $\tau' = 3.16$ and $\eta_0 = 0.59$ for optimally spaced stimuli and then examine eight locations around the 90% contour corresponding to this optimal point. Results for $\mu(t')$ and $\eta'(t')$ are shown in the upper and lower graphs of each panel. Each panel also schematically denotes the contour, with the solid circle indicating

the location on the contour used in that panel (or the central, optimal point). Panel E gives the optimal point and the others the eight points around the 90% contour. We show $\mu(t')$ for both the spaced (solid lines) and massed (dashed lines) stimuli, and we also show the rectified, relative difference between $\mu(t')$ for the spaced stimulus and the maximum value of $\mu(t')$ for the massed stimulus (dotted lines).

Figure 11: Realisations of $h(t')$ compared to the mean signal $\mu(t')$. For $\Theta = 8$ and $\rho = 6$ we select the optimal parameter $\tau' = 3.16$ and $\eta_0 = 0.59$ and corresponding optimally spaced times $t'_i \in \{0, 21, 37, 51, 64, 77, 90\}$. We plot $\mu(t')$ (solid line) and four completely representative examples (dotted lines) of particular realisations of $h(t')$, as functions of t'. Also shown in grey are the 1σ and 2σ regions around the mean. The mean does not depend on the number of synapses, N, but $\sigma^2(t')$ does, so for the particular realisations we set N as indicated in each panel.

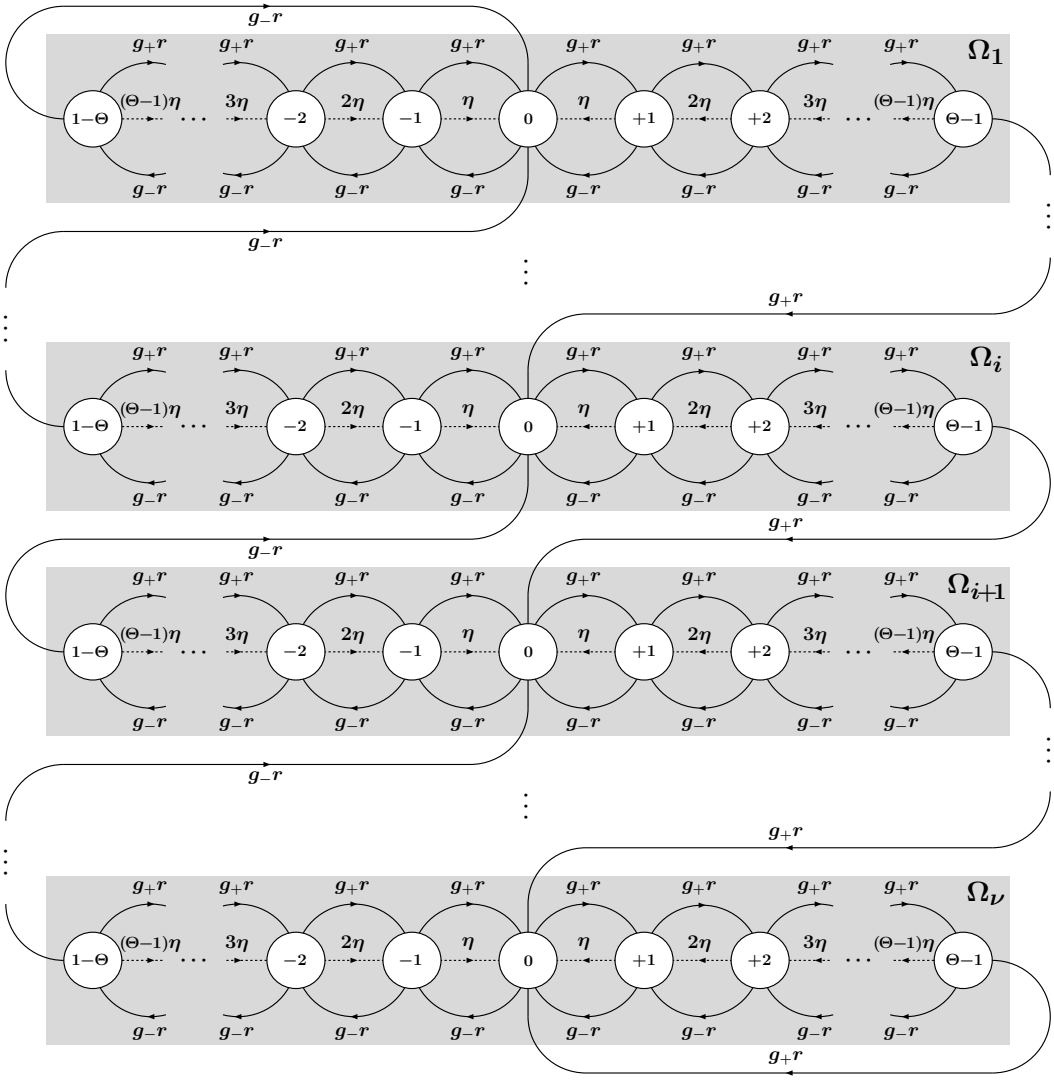
Ω_1
Ω_i
Ω_{i+1}
Ω_ν
g_-r
g_+r
$1-\Theta$
$(\Theta-1)\eta$
3η
-2
2η
-1
η
0
η
$+1$
2η
$+2$
3η
$(\Theta-1)\eta$
$\Theta-1$

Figure 1

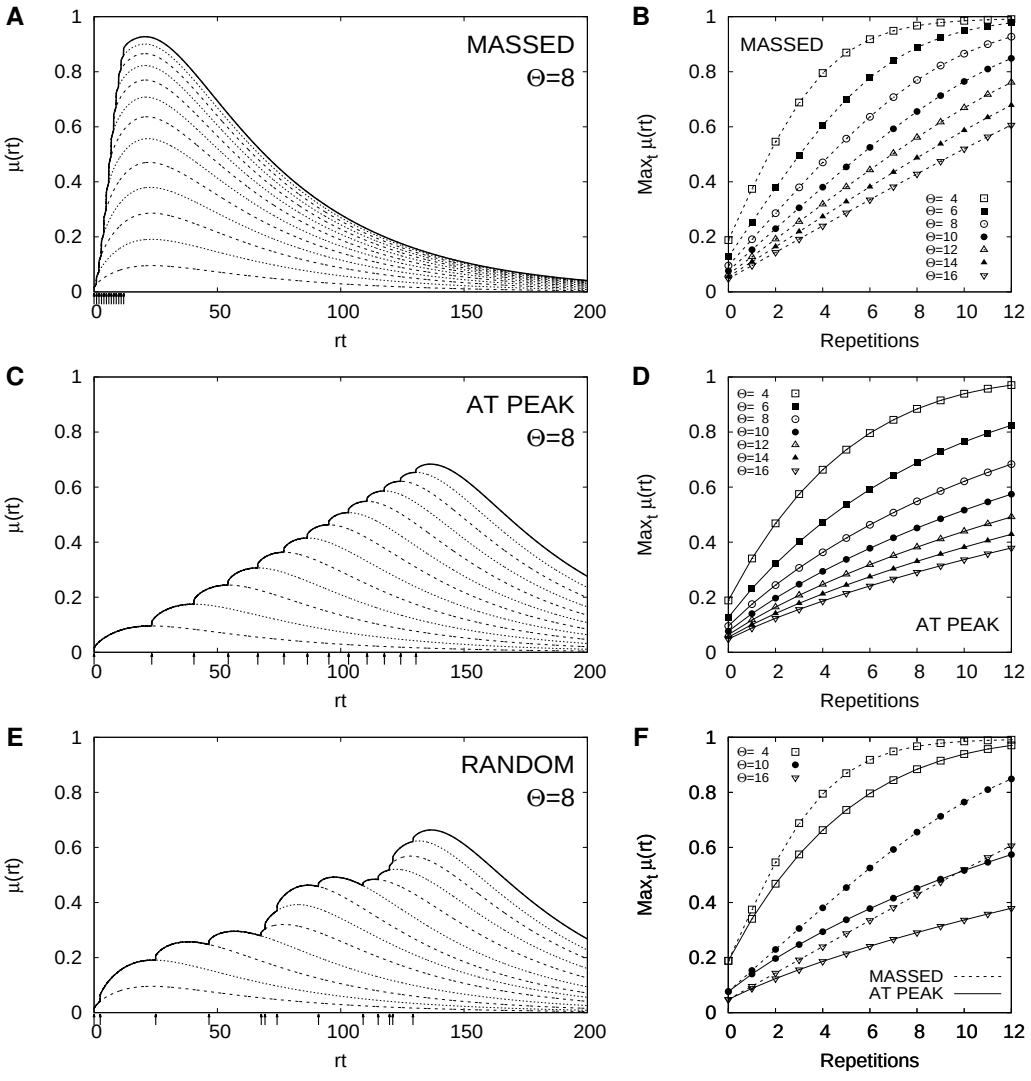
A
MASSED
Θ=8
μ(rt)
rt
B
MASSED
Max_t μ(rt)
Repetitions
Θ= 4
Θ= 6
Θ= 8
Θ=10
Θ=12
Θ=14
Θ=16
C
AT PEAK
Θ=8
μ(rt)
rt
D
Max_t μ(rt)
Repetitions
AT PEAK
E
RANDOM
Θ=8
μ(rt)
rt
F
Max_t μ(rt)
Repetitions
MASSED
AT PEAK

Figure 2

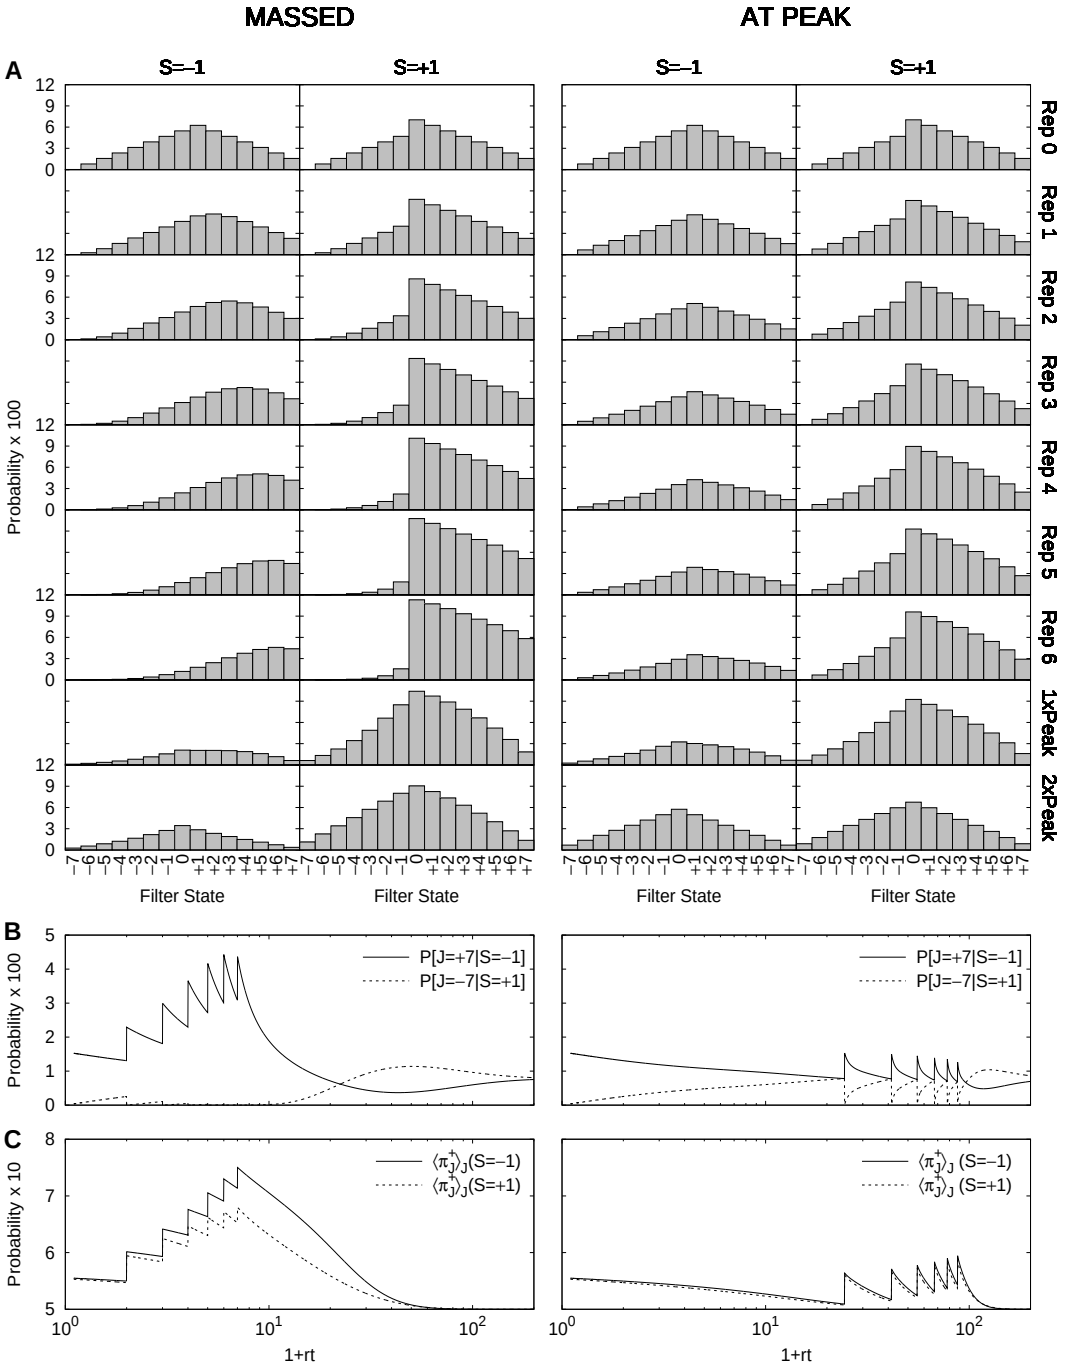
MASSED
AT PEAK
A
S=–1
S=+1
Rep 0
Rep 1
Rep 2
Rep 3
Rep 4
Rep 5
Rep 6
1xPeak
2xPeak
Probability x 100
Filter State
B
P[J=+7|S=–1]
P[J=–7|S=+1]
C
1+rt

Figure 3

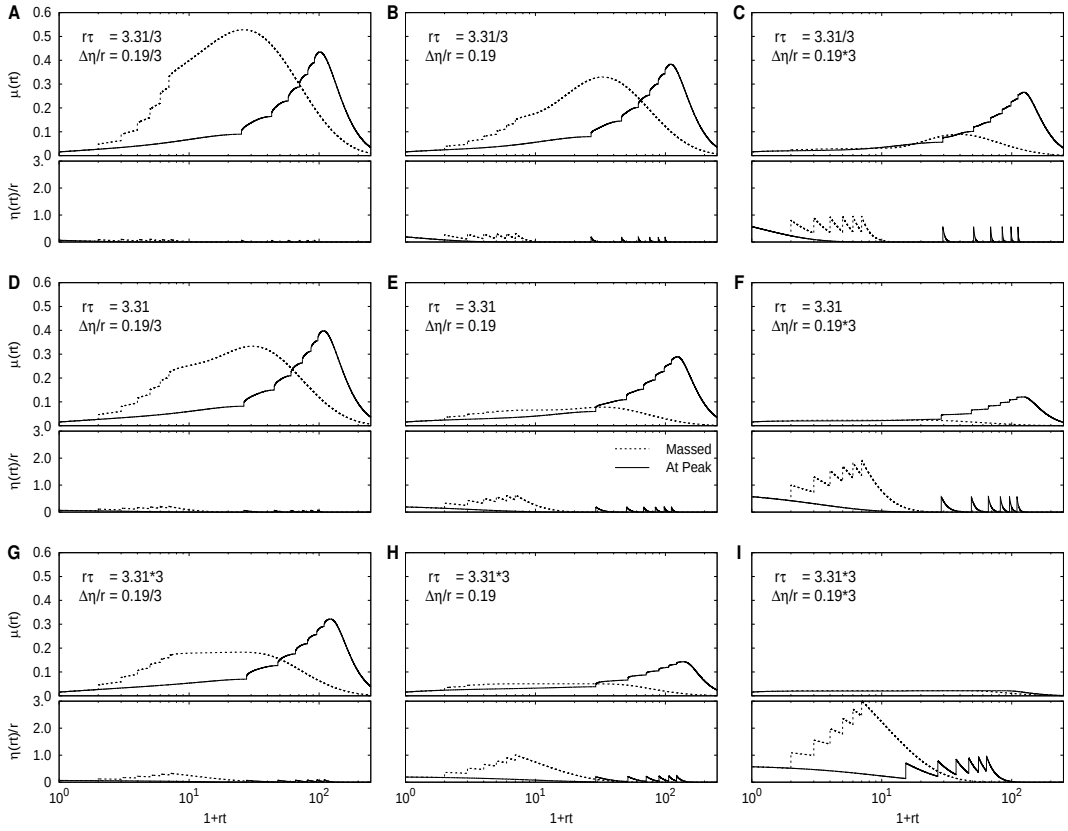
A
B
C
D
E
F
G
H
I
rτ = 3.31/3
Δη/r = 0.19/3
rτ = 3.31/3
Δη/r = 0.19
rτ = 3.31/3
Δη/r = 0.19*3
rτ = 3.31
Δη/r = 0.19/3
rτ = 3.31
Δη/r = 0.19
rτ = 3.31
Δη/r = 0.19*3
rτ = 3.31*3
Δη/r = 0.19/3
rτ = 3.31*3
Δη/r = 0.19
rτ = 3.31*3
Δη/r = 0.19*3
μ(rt)
η(rt)/r
Massed
At Peak
1+rt

Figure 4

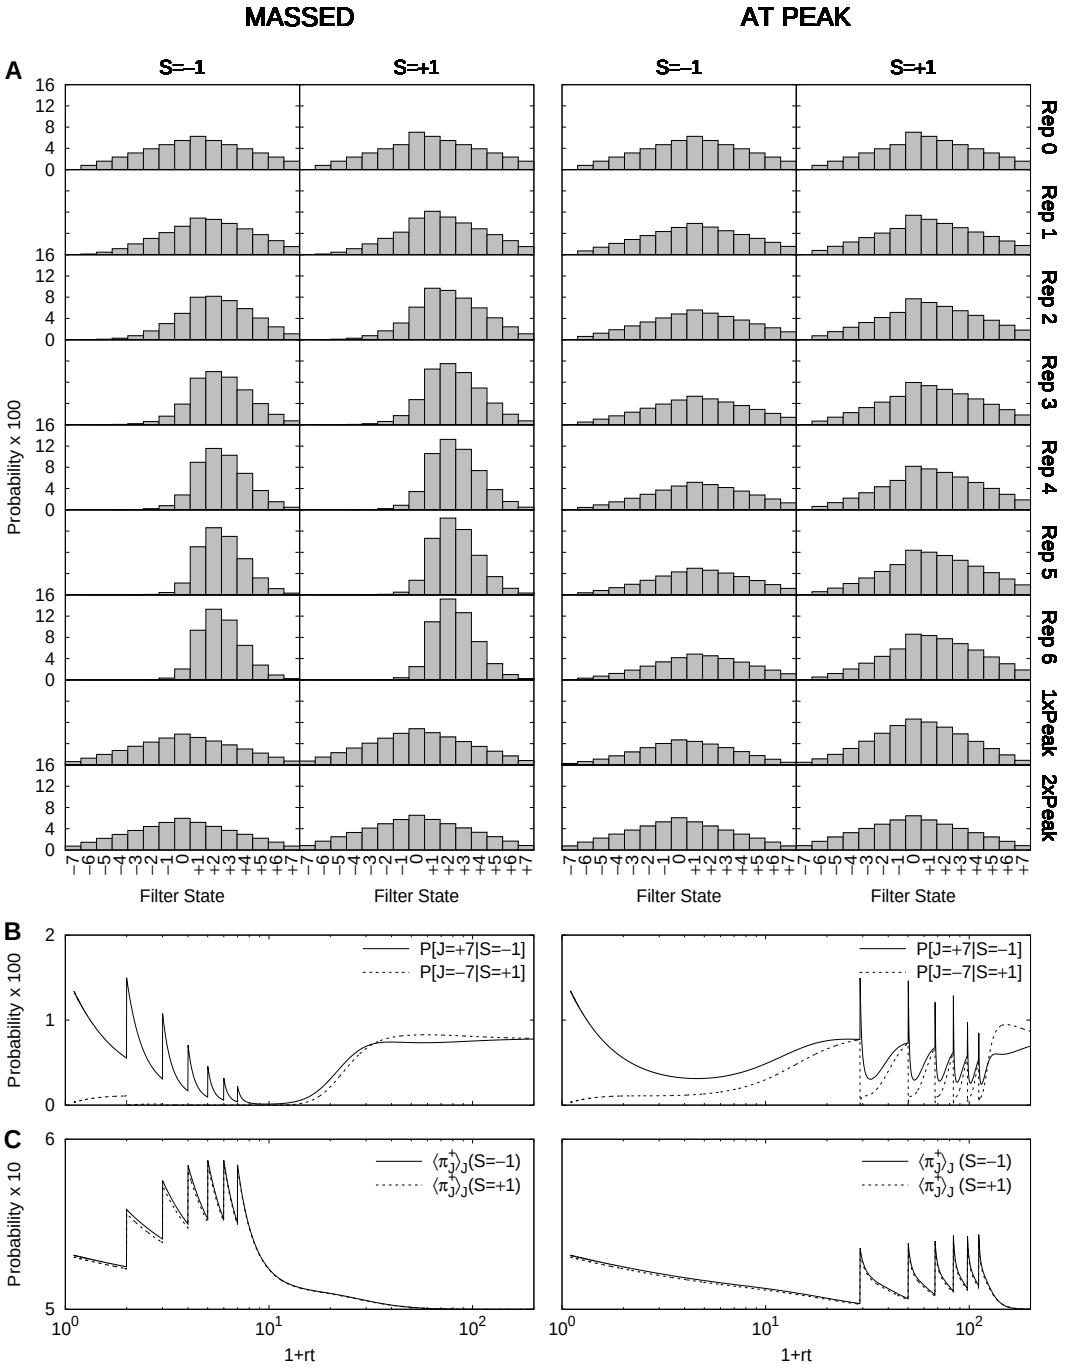
MASSED
AT PEAK
A
S=–1
S=+1
Rep 0
Rep 1
Rep 2
Rep 3
Rep 4
Rep 5
Rep 6
1xPeak
2xPeak
Probability x 100
Filter State
B
P[J=+7|S=–1]
P[J=–7|S=+1]
C
1+rt

Figure 5

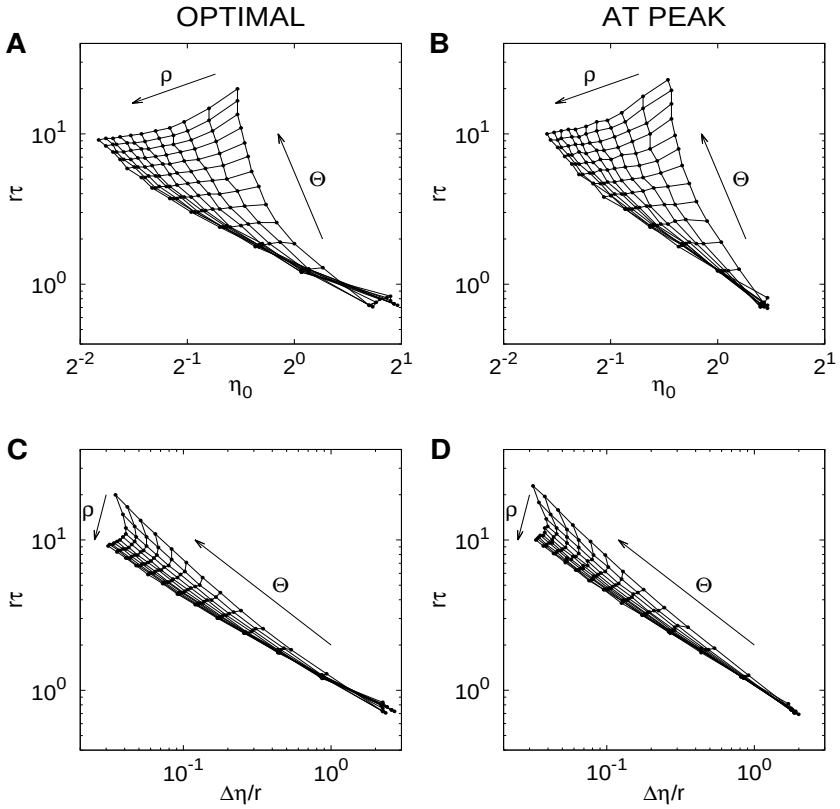
OPTIMAL
AT PEAK
A
B
C
D
ρ
Θ
r
10^1
10^0
2^{-2}
2^{-1}
2^0
2^1
η_0
10^{-1}
10^0
$\Delta\eta/r$

Figure 6

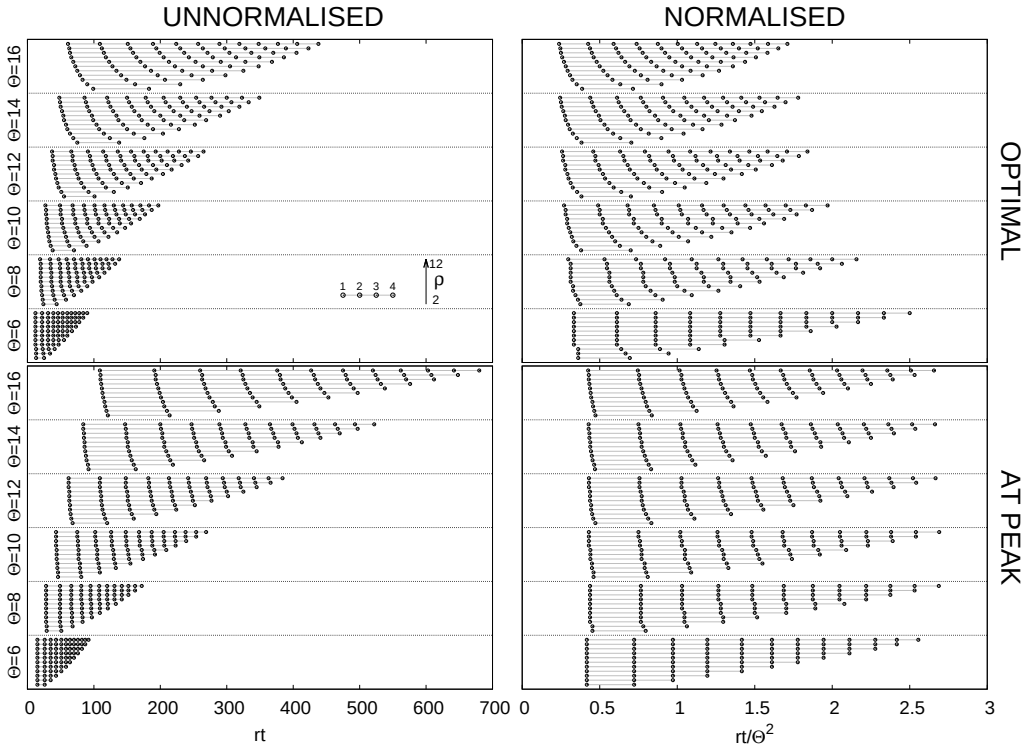
UNNORMALISED
NORMALISED
OPTIMAL
AT PEAK
Θ=6
Θ=8
Θ=10
Θ=12
Θ=14
Θ=16
1 2 3 4
ρ
12
2
0
100
200
300
400
500
600
700
rt
0
0.5
1
1.5
2
2.5
3
rt/Θ²

Figure 7

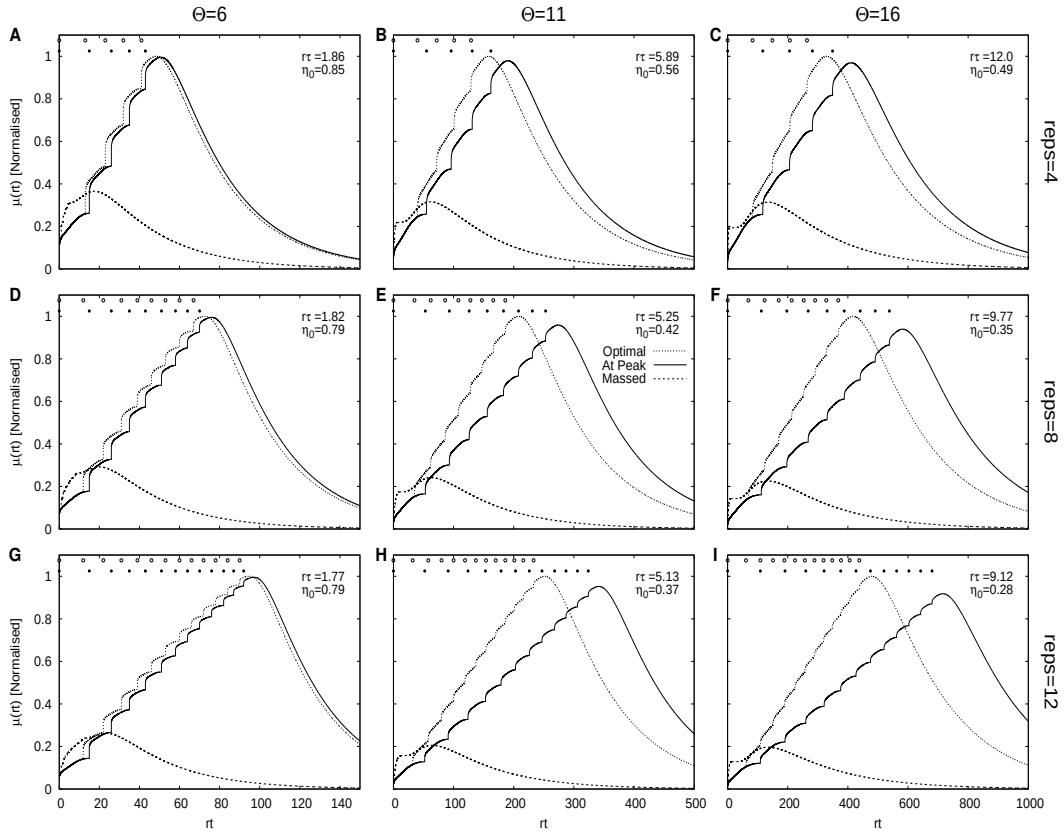
Θ=6
Θ=11
Θ=16
A
B
C
D
E
F
G
H
I
rτ =1.86
η0=0.85
rτ =5.89
η0=0.56
rτ =12.0
η0=0.49
rτ =1.82
η0=0.79
rτ =5.25
η0=0.42
rτ =9.77
η0=0.35
rτ =1.77
η0=0.79
rτ =5.13
η0=0.37
rτ =9.12
η0=0.28
Optimal
At Peak
Massed
μ(rt) [Normalised]
rt
reps=4
reps=8
reps=12

Figure 8

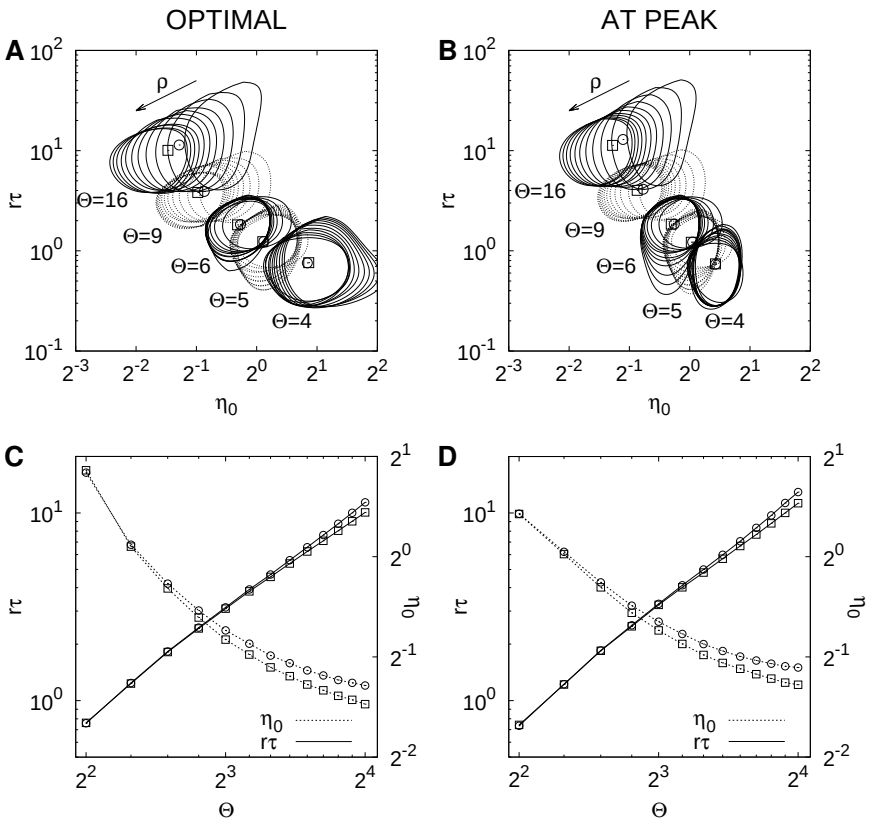

Figure 9

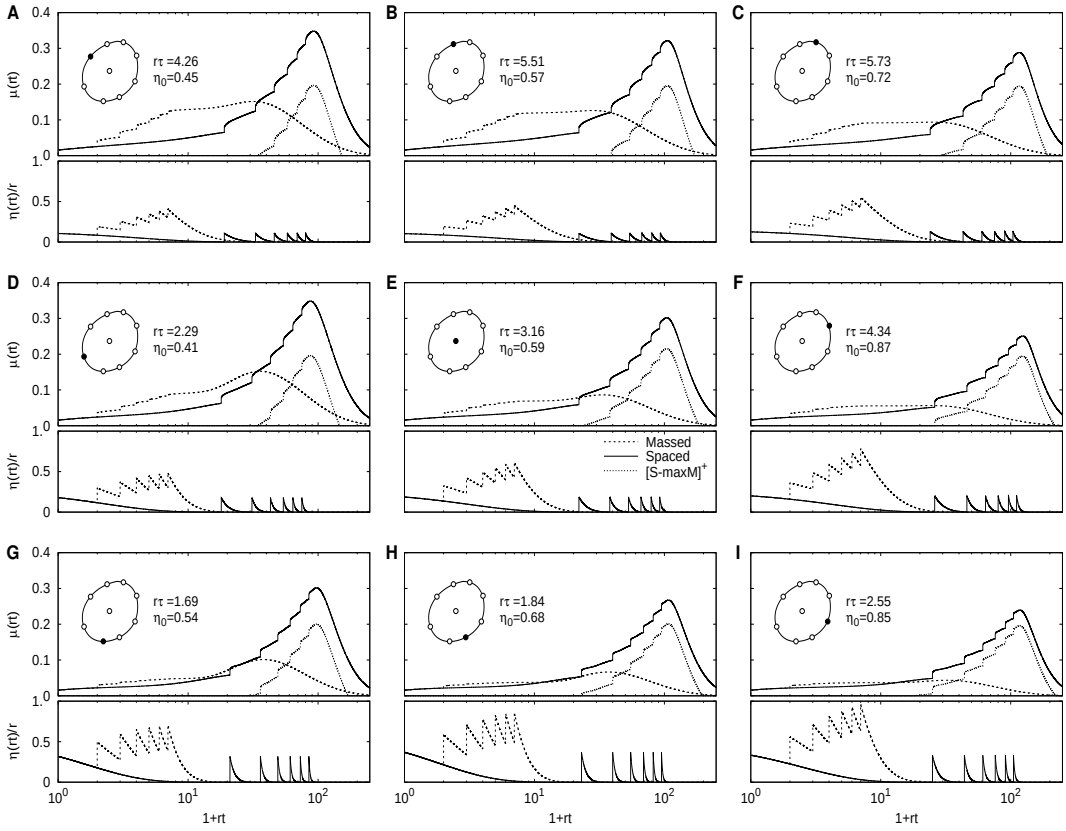
A
B
C
D
E
F
G
H
I
rτ =4.26
η0=0.45
rτ =5.51
η0=0.57
rτ =5.73
η0=0.72
rτ =2.29
η0=0.41
rτ =3.16
η0=0.59
rτ =4.34
η0=0.87
rτ =1.69
η0=0.54
rτ =1.84
η0=0.68
rτ =2.55
η0=0.85
μ(rt)
η(rt)/r
Massed
Spaced
[S-maxM]+
1+rt

Figure 10

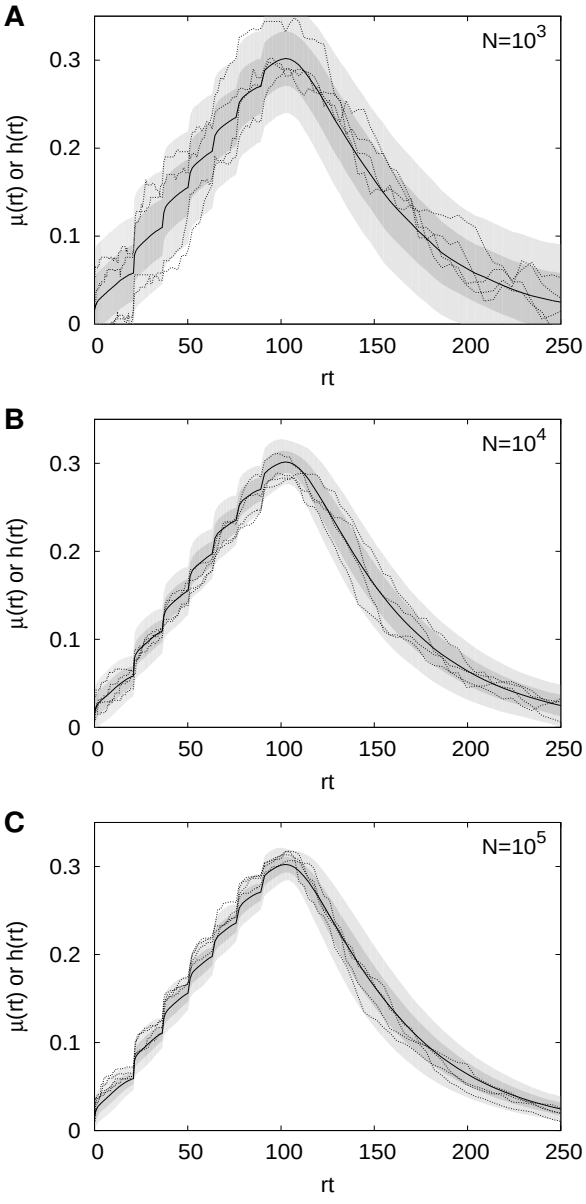

Figure 11