LETTER Communicated by Walter Senn

Stable Competitive Dynamics Emerge from Multispike Interactions in a Stochastic Model of Spike-Timing-Dependent Plasticity

Peter A. Appleby
p.appleby@biologie.hu-berlin.de
Institute for Theoretical Biology, Humboldt-Universität zu Berlin, Invalidenstrasse 43, D-10115 Berlin, Germany

Terry Elliott
te@ecs.soton.ac.uk
Department of Electronics and Computer Science, University of Southampton, Highfield, Southampton SO17 1BJ, U.K.

In earlier work we presented a stochastic model of spike-timing-dependent plasticity (STDP) in which STDP emerges only at the level of temporal or spatial synaptic ensembles. We derived the two-spike interaction function from this model and showed that it exhibits an STDP-like form. Here, we extend this work by examining the general n-spike interaction functions that may be derived from the model. A comparison between the two-spike interaction function and the higher-order interaction functions reveals profound differences. In particular, we show that the two-spike interaction function cannot support stable, competitive synaptic plasticity, such as that seen during neuronal development, without including modifications designed specifically to stabilize its behavior. In contrast, we show that all the higher-order interaction functions exhibit a fixed-point structure consistent with the presence of competitive synaptic dynamics. This difference originates in the unification of our proposed "switch" mechanism for synaptic plasticity, coupling synaptic depression and synaptic potentiation processes together. While three or more spikes are required to probe this coupling, two spikes can never do so. We conclude that this coupling is critical to the presence of competitive dynamics and that multispike interactions are therefore vital to understanding synaptic competition.

1 Introduction

Spike-timing-dependent plasticity (STDP) has become a topic of much interest. In STDP, the exact timing of pre- and postsynaptic stimulation determines the degree and polarity of change in synaptic strength (for review, see Roberts & Bell, 2002). This is in contrast to conventional, rate-based

 © 2006 Massachusetts Institute of Technology

synaptic plasticity (Bliss & Lømo, 1973; Gustafsson, Wigström, Abraham, & Huang, 1987; Dudek & Bear, 1992), where the governing factor is the rate of pre- and postsynaptic firing. STDP is apparently widespread (Bi & Poo, 1998; Zhang, Tao, Holt, Harris, & Poo, 1998; Froemke & Dan, 2002), and as a result, the learning properties of STDP rules are of potentially far-reaching consequence. In theoretical studies, STDP rules have been explored in a variety of phenomenological (Song, Miller, & Abbott, 2000; van Rossum, Bi, & Turrigiano, 2000; Izhikevich & Desai, 2003) and biophysical (Castellani, Quinlan, Cooper, & Shouval, 2001; Karmarkar & Buonomano, 2002; Shouval, Bear, & Cooper, 2002) models. In all cases, the basic phenomenology of STDP can be reproduced, but the models are often sensitive to the choice of parameters or difficult to generalize to multispike interactions (Froemke & Dan, 2002).

A key feature of developmental synaptic plasticity in the visual system and elsewhere is that it is activity dependent and competitive in character, with stable, segregated patterns of afferent innervation emerging as development proceeds (Purves & Lichtman, 1985). Although rate-based synaptic plasticity rules have been generally successful in reproducing much of the phenomenology of developmental plasticity (see, e.g., van Ooyen, 2003), it is natural to wonder whether models of STDP can also account for such processes, especially given that STDP is expressed, for example, in the developing visual system (Zhang et al., 1998; Schuett, Bonhoeffer, & Hübener, 2001) and that STDP may in fact underlie rate-based synaptic plasticity. If STDP really does underlie rate-based synaptic plasticity, then the connection between STDP models and rate-based plasticity models must be established, and the capacity of STDP models to exhibit the stable, competitive dynamics characteristic of developmental synaptic plasticity must also be demonstrated. To be sure, models of STDP have been related to one particular rate-based plasticity model, the Bienenstock-Cooper-Munro (BCM) (Bienenstock, Cooper, & Munro, 1982) model (Izhikevich & Desai, 2003). Furthermore, a variety of bimodal (Song et al., 2000) and unimodal (van Rossum et al., 2000) distributions of synaptic weights can be obtained under simple STDP rules in the presence of additional constraints. However, the conditions under which STDP can in general support stable, competitive dynamics, and the relationship between rate- and timing-based models of plasticity more generally, have only been partially explored.

Existing models of STDP typically assume that STDP operates at each individual synapse. This places a heavy computational burden on the synapse, requiring an array of spike coincidence detection machinery to be present that can represent spike time differences to millisecond accuracies and adjust synaptic weights in a continuous, graded fashion. As experimental work on STDP typically involves measuring the plasticity of multisynapse connections across multiple spike pairings, the observed STDP plasticity rule could, however, emerge from the averaging of some simpler synaptic rule. In earlier work, we postulated such a model of STDP,

in which something akin to a three-state switch governs individual synaptic changes (Appleby & Elliott, 2005), with changes in synaptic strength occurring in discrete jumps of fixed magnitude (cf. Peterson, Malenka, Nicoll, & Hopfield, 1998; O'Connor, Wittenberg, & Wang, 2005). The overall change in strength in response to repeated spike pairings was then shown to be of an STDP-like form. The observed STDP rule can therefore emerge as a temporal and spatial average over multiple synapses and multiple spike pairings, with synapses required only to perform a minimal level of coincidence detection. In addition, an explanation of spike triplet interactions (Froemke & Dan, 2002) emerged as a natural consequence of the structure of the switch rule. Using simple probability theory to consider the interaction of any two spikes in a spike train, we derived an averaged, rate-based learning rule and derived a set of constraints under which a BCM-like form may be obtained.

Elsewhere (Appleby & Elliott, 2006), we perform a fuller analytical treatment that takes into account the higher-order spike interactions that occur when examining spike trains composed of more than two spikes, and derive a generating function for the general n-spike, averaged learning rule that emerges from the model. For convenience, we refer to interactions involving three or more spikes as *multispike interactions*. In this article, we examine the learning dynamics of multispike rules in much greater detail, with the aim of exploring the connections between STDP and rate-based plasticity in our model. Extending the simple probabilistic argument presented in Appleby & Elliott (2005), we present averaged, rate-based rules for cases where spike interactions are limited to two, three, or four spikes and the case where interactions are not limited in any way. We observe that the two-spike learning rule, despite possessing a promisingly BCM-like form, leads inevitably to pathological learning behaviors unless further modifications are made to stabilize its dynamics. We then explore the effect of allowing spike interactions to extend beyond two spikes and find that unlike the two-spike rule, all multispike learning rules can exhibit stable, competitive dynamics without introducing any additional modifications. We explore analytically the fixed-point structure of the switch rule under the assumption that afferents stochastically vary about the same mean firing rate and find that the presence of higher-order corrections in multispike rules changes the fixed-point structure compared to the two-spike rule. We argue that this change occurs because potentiation and depression are coupled in the switch model in the presence of multispike interactions. We demonstrate, numerically, that the multispike learning rules are computational modes available to a real neuron and find a constraint on the magnitude of plasticity below which the rate-based behavior becomes dominant. Finally, we perform several numerical simulations, including a large-scale simulation of ocular dominance column (ODC) development, that demonstrates that the observed fixed-point structure leads to stable, competitive dynamics in a large-scale system.

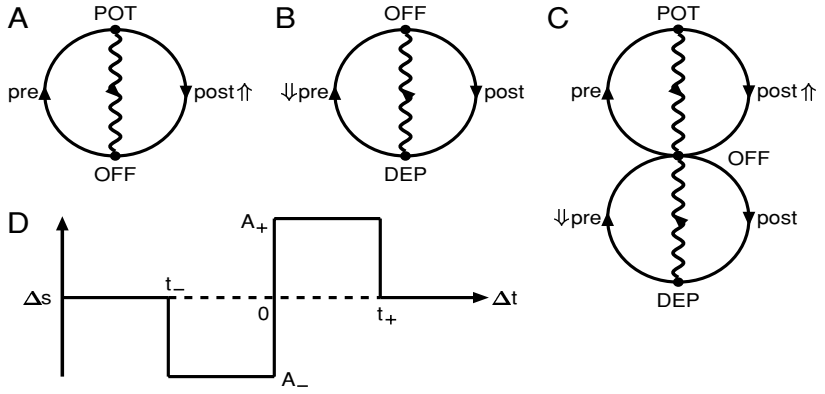

Figure 1: A schematic representation of the switch and its resulting plasticity rule. (A) A switch to account for potentiation induced by a presynaptic spike preceding a postsynaptic spike. (B) An equivalent switch to account for depression induced when the spike order is reversed. Both switches can be unified into a single switch (C) with a common *OFF* state. Such a switch cannot simultaneously exist in both the *DEP* and *POT* states, but only in one state at a time. (D) shows the change in synaptic strength evoked under the unified three-state switch for a representative spike pair at various spike timings. Wavy lines represent stochastic transitions to the *OFF* state without any induced potentiation or depression. Semicircles represent transitions induced by a pre- or postsynaptic spike. The arrows ⇑ and ⇓ indicate the induction of potentiation and depression, respectively.

2 Summary of Switch Model

We first give a brief overview of our synaptic switch model presented in Appleby & Elliott (2005). We propose that a pair of spikes, one pre- and one postsynaptic, will potentiate or depress a given synapse by a fixed amount, $A_{\pm}$, subject only to the requirement that the second spike occurs within a finite time window relative to the first spike. Outside this time window, the second spike does not evoke any change in synaptic strength. The duration of this time window is not fixed but is taken to be a stochastic quantity governed by some probability distribution. The degree of potentiation or depression does not depend on the time difference between events. Rather, we require only that the second event occurs within a finite time window after the first event for the fixed steps, $A_{\pm}$, to occur.

This simple modification rule could be embodied by some biological, synaptic switch mechanism similar to that shown in Figure 1A, which represents the process leading to potentiation under a pre- and then postsynaptic spike pairing. The synapse initially resides in a resting state that we label the *OFF* state. A presynaptic spike elevates the synapse into a different

functional state, which we label the POT state. A postsynaptic spike occurring while the switch is in the POT state immediately returns the switch to the *OFF* state and induces an associated potentiation of synaptic strength of magnitude A_+. In the absence of a postsynaptic spike, the switch will move from the *POT* state back to the *OFF* state in a stochastic manner, governed by some probability distribution. Such a stochastic transition does not induce any change in synaptic strength. We postulate an identical switch to account for depression (see Figure 1B), which operates in the same manner. This time, an initial postsynaptic spike activates the synapse, moving it into a third functional state, which we label the *DEP* state. A presynaptic spike occurring while the switch is in the *DEP* state immediately returns the switch to the *OFF* state and induces an associated depression of synaptic strength of magnitude A_-. In the absence of a presynaptic spike, the switch will spontaneously move from the *DEP* state back to the *OFF* state, again in a stochastic manner, without any change in synaptic strength.

We choose to unify these two switches into a single three-state switch (see Figure 1C). Although any number of additional states and transitions may be postulated, we find that this simple switch mechanism is sufficient to reproduce a variety of STDP results (Appleby & Elliott, 2005). Although the unification of the switch was initially motivated on the grounds of simplicity, unifying the switch also has the effect of coupling potentiation and depression together. This coupling takes place in the sense that once the synapse enters the *POT* state, it cannot subsequently enter the *DEP* state without first returning to the *OFF* state (and vice versa). We will show later that this coupling of potentiation and depression dramatically affects the dynamical landscape of the multispike learning rules derivable from the switch model compared to two uncoupled switches. A depiction of the step change in synaptic strength that a spike pair induces as a function of spike timing is shown in Figure 1D. The magnitude of change is fixed: there is no dependence on the time difference between pre- and postsynaptic events, and synapses undergo an all-or-none potentiation or depression.

We distinguish two forms of the model based on its response to additional presynaptic spiking once the synapse is already in the *POT* state. In the *resetting model*, additional presynaptic spiking while in the *POT* state resets the stochastic process that governs the transition back to *OFF*. In the *non-resetting model*, additional presynaptic spiking has no further effect. The behavior of the depressing lobe of the switch is taken to be the same as the potentiating lobe. In the resetting model, therefore, postsynaptic spiking is capable of resetting the stochastic process governing the $DEP \to OFF$ transition in an identical manner to that of presynaptic spiking resetting the stochastic $POT \to OFF$ process. In the resetting form of a multispike rule, there is a tendency to experience more potentiation or depression than in the nonresetting form, as the effect of resetting is to increase the proportion of time that the synapse spends in the *POT* or *DEP* states compared to the *OFF* state. Resetting also has the effect of discarding the spike history. Once

elevated to the *POT* state, each subsequent presynaptic spike resets the stochastic process and eradicates all trace of its predecessors. This gives rise to a lack of memory in the sense that if the synapse is in the *POT* state, then only the time since the last presynaptic spike, not the times of all presynaptic spikes since the initial elevation to the *POT* state, is relevant. A synapse that has been reset n times by n additional presynaptic spikes is therefore identical to a synapse that has just been elevated to the *POT* state.

In the experimental STDP protocols that we consider (Bi & Poo, 1998), we are limited to precisely two spikes, which, due to the structure of the switch, will never be enough to probe the resetting behavior. The choice of resetting or nonresetting is arbitrary to some extent, as the basic phenomenology of STDP is reproduced under both forms of the model. It is only when considering spike trains comprising three or more spikes that the choice of whether to allow resetting begins to make an impact. This impact is, however, minimal, in the sense that the resetting and nonresetting forms display qualitatively similar learning behavior. The motivation for considering resetting lies in the observation that the resetting form is particularly convenient for deriving the spike interaction function for any number of spikes (Appleby & Elliott, 2006).

We assume, as before, that an afferent makes multiple synapses onto a target cell and that the overall connection strength between the two cells is the linear sum of each individual synaptic strength. The synapses are treated independently, which, due to the stochastic nature of the synaptic modification rule, means that the synapses comprising a connection will often be in different states. It is therefore the spatial average over synapses, and the temporal average over spike pairs, that determines the overall change in connection strength. This overall change in connection strength for a multisynapse connection as a function of spike timing is qualitatively STDP-like (Appleby & Elliott, 2005). Thus, when the changes are viewed at the level of multisynapse connections, the averaging of the proposed synaptic switch rule leads directly to an STDP-like modification curve.

2.1 Multispike Interaction Functions. We may calculate the expected response of a synapse to a train of spikes using simple probability theory. We first reproduce the derivation of the two-spike interaction function (Appleby & Elliott, 2005) and then state the three-, four-, and ∞-spike rules. For notational convenience, we denote a presynaptic spike by the symbol π and a postsynaptic spike by the symbol p. Pre- and postsynaptic firing are assumed to be independent Poisson processes with rates λ_π and λ_p, respectively, and we set $\beta = \lambda_\pi + \lambda_p$. Because they are independent, the combined pre- and postsynaptic spike sequences form a single Poisson process of overall rate β. The probability density function for the interspike interval of this joint process is $f_T(t) = \beta \exp(-\beta t)$, where t denotes time. The probability that any particular spike in this combined train is a π spike is $\lambda'_\pi = \lambda_\pi/\beta$, while the probability of its being a p spike is $\lambda'_p = \lambda_p/\beta$. We

use the subscript + to refer to quantities governing the potentiating lobe of the switch, and the subscript − to refer to their counterparts governing the depressing lobe of the switch. For example, the magnitudes of plasticity are denoted A_+ for potentiation and A_- for depression, respectively. The potentiating and depressing lobes of the switch do not necessarily share the same parameters, representing the possibility that the two processes have at least some degree of independence. We assume that the probability density functions $f_\pm(t)$ for the stochastic transitions $POT \to OFF$ and $DEP \to OFF$ (see Figures 1A and 1B) are given by integer-order gamma distributions, with probability density functions

$$f_\pm(t) = \frac{(t/\tau_\pm)^{n_\pm - 1}}{(n_\pm - 1)!}\frac{1}{\tau_\pm}\exp(-t/\tau_\pm), \tag{2.1}$$

where $n_\pm$ are the integer orders and $\tau_\pm$ are the characteristic timescales associated with these stochastic processes. For $n_\pm = 1$, these are just simple exponential distributions. The lack of memory property associated with such exponential distributions means that only the time since the most recent spike is relevant. For $n_\pm = 1$, the choice of resetting or nonresetting is therefore irrelevant, as the exponential stochastic processes themselves erase the spike history.

We may calculate the expected response to a typical spike train comprising n spikes. This involves deconstructing the spike train into all the different possible spike combinations, then calculating the contribution from each combination to the expected change in synaptic strength. Because we have two different spike types (pre- and postsynaptic), we have a total of 2^n possible spike trains of length n. Integrating out the interspike intervals and averaging across all possible spike combinations gives an unconditional expectation value for the average change induced by a typical n-spike train.

Consider first a two-spike train. The spike train may manifest itself in one of four possible spike sequences, $\pi\pi$, πp, $p\pi$ or pp, occurring with probabilities $p_{\sigma_1\sigma_2} = \lambda'_{\sigma_1}\lambda'_{\sigma_2}$, where $\sigma_1, \sigma_2 \in \{\pi, p\}$. Each of these spike sequences gives a contribution, denoted by $\Delta S_{\sigma_1\sigma_2}$, to the expected change for the typical two-spike train, denoted by ΔS_2. This contribution is obtained by conditioning on the state of the synapse when the second spike arrives and calculating a mean expected change. Consider the spike sequence πp. The mean change in synaptic efficacy triggered by a πp spike pair of time difference t is the amplitude of synaptic plasticity, A_+, multiplied by the probability that the switch is still *ON* at time t after the first spike, $P^+(t)$. Thus, the conditional expectation value for the change in synaptic efficacy, $\Delta S_{\pi p}(t)$, given a πp spike time interval t, is just

$$\Delta S_{\pi p}(t) = +A_+ P^+(t). \tag{2.2}$$

Similarly, an identical argument gives

$$\Delta S_{p\pi}(t) = -A_- P^-(t), \tag{2.3}$$

for the sequence $p\pi$, where $P^{\pm}(t)$ are the probabilities that the synapse is still in the elevated state when the second spike arrives. These probabilities are just

$$P^{\pm}(t) = \int_t^{\infty} dt' f_{\pm}(t'), \tag{2.4}$$

which, for the gamma distributions in equation 2.1, integrate to give

$$P^{\pm}(t) = e_{n_{\pm}}(t/\tau_{\pm}) \exp(-t/\tau_{\pm}), \tag{2.5}$$

where $e_n(x) = \sum_{i=0}^{n-1} x^i / i!$ is the truncated exponential series. The spike patterns $\pi\pi$ and pp cannot cause a change in synaptic strength under our switch rule, so $\Delta S_{\pi\pi}(t) = \Delta S_{pp}(t) = 0$. Averaging across all four possible manifestations of the two-spike train gives the expected change in synaptic strength for a typical two-spike train. This expression is a conditional expectation value due to the dependence on the interspike interval between the pre- and postsynaptic spikes. Integrating out the interspike time intervals according to their exponential distributions, $f_T(t)$, gives the unconditional expectation value for the change in synaptic strength induced by a two-spike train. The two-spike rule may therefore be written as

$$\Delta S_2 = \sum_{\sigma_1,\sigma_2 \in \{\pi,p\}} p_{\sigma_1\sigma_2} \int_0^{\infty} dt_0 \int_0^{\infty} dt_1 f_T(t_0) f_T(t_1) \Delta S_{\sigma_1\sigma_2}(t_1). \tag{2.6}$$

Defining

$$K_1^{\pm}(\beta) = \int_0^{\infty} dt f_T(t) P^{\pm}(t) = 1 - \frac{1}{(1+\tau_{\pm}\beta)^{n_{\pm}}} \tag{2.7}$$

and $\tilde{K}_1^+(\beta) = \lambda'_\pi K_1^+(\beta)$ and $\tilde{K}_1^-(\beta) = \lambda'_p K_1^-(\beta)$, it is straightforward to show that the two-spike rule takes the form

$$\Delta S_2 = \lambda'_p A_+ \tilde{K}_1^+(\beta) - \lambda'_\pi A_- \tilde{K}_1^-(\beta). \tag{2.8}$$

This equation is an analytical expression for the expected change in synaptic efficacy induced by a two-spike train at given pre- and postsynaptic firing rates, λ_π and λ_p. As a two-spike train can never probe the resetting properties of the switch, this form is valid for both resetting and nonresetting models.

For an n-spike train, equation 2.6 becomes

$$\Delta S_n = \sum_{\sigma_1,\ldots,\sigma_n \in \{\pi,p\}} p_{\sigma_1\cdots\sigma_n} \int_0^\infty dt_0 \cdots dt_{n-1}\; f_T(t_0)\cdots f_T(t_{n-1}) \times \Delta S_{\sigma_1\cdots\sigma_n}(t_1,\ldots,t_{n-1}), \tag{2.9}$$

where $\Delta S_{\sigma_1\cdots\sigma_n}(t_1,\ldots,t_{n-1})$ is the expected change induced by the spike sequence $\sigma_1\ldots\sigma_n$ at interspike intervals $t_1,\ldots,t_{n-1}$, and $p_{\sigma_1,\ldots,\sigma_n} = \lambda'_{\sigma_1}\times\ldots\times\lambda'_{\sigma_n}$. The function $\Delta S_{\sigma_1,\ldots,\sigma_n}$ does not depend on the time of the first spike, t_0, only on the interspike intervals. It may be calculated directly by conditioning on the state of the synapse as each spike in the sequence arrives. Although cumbersome, this method allows ΔS_n to be explicitly evaluated for small n. For larger n, alternative methods based on formulating a recurrence relation in the ΔS_n can be used (Appleby & Elliott, 2006). In the nonresetting model, we find by direct calculation that

$$\begin{aligned}\Delta S_3^{NR} = {} & \lambda'_p A_+ \left[F_0\tilde{K}_2^+(\beta) + (F_0+F_1)\tilde{K}_1^+(\beta)\right] \\ & - \lambda'_\pi A_- \left[F_0\tilde{K}_2^-(\beta) + (F_0+F_1)\tilde{K}_1^-(\beta)\right],\end{aligned} \tag{2.10}$$

and

$$\begin{aligned}\Delta S_4^{NR} = {} & \lambda'_p A_+ \left[F_0\tilde{K}_3^+(\beta) + (F_0+F_1)\tilde{K}_2^+(\beta) + (F_0+F_1+F_2)\tilde{K}_1^+(\beta)\right] \\ & - \lambda'_\pi A_- \left[F_0\tilde{K}_3^-(\beta) + (F_0+F_1)\tilde{K}_2^-(\beta) + (F_0+F_1+F_2)\tilde{K}_1^-(\beta)\right],\end{aligned} \tag{2.11}$$

where

$$\begin{aligned}K_l^\pm(\beta) &= \int_0^\infty dt_1\cdots dt_l\, f_T(t_1)\cdots f_T(t_l) P^\pm(t_1+\cdots+t_l) \\ &= \left(\frac{\tau_\pm\beta}{1+\tau_\pm\beta}\right)^l \sum_{i=0}^{n_\pm-1} \frac{(i+l-1)!}{i!(l-1)!}\frac{1}{(1+\tau_\pm\beta)^i},\end{aligned} \tag{2.12}$$

and

$$\tilde{K}_l^+(\beta) = {\lambda'_\pi}^l K_l^+(\beta), \tag{2.13}$$

$$\tilde{K}_l^-(\beta) = {\lambda'_p}^l K_l^-(\beta), \tag{2.14}$$

and

$$F_0 = 1, \tag{2.15}$$

$$F_1 = 1 - (\tilde{K}_1^+(\beta) + \tilde{K}_1^-(\beta)), \tag{2.16}$$

$$F_2 = 1 - (\tilde{K}_1^+(\beta) + \tilde{K}_1^-(\beta)) - \left[\tilde{K}_2^+(\beta) + \tilde{K}_2^-(\beta) - (\tilde{K}_1^+(\beta) + \tilde{K}_1^-(\beta))^2\right]. \tag{2.17}$$

The resetting model replaces $P^{\pm}(t_1 + \cdots + t_l)$ in the definition of $K_l^{\pm}(\beta)$ by the product $P^{\pm}(t_1) \times \cdots \times P^{\pm}(t_l)$. Hence, for the resetting model, we have that

$$K_l^{\pm}(\beta) = \left[\int_0^{\infty} dt f_T(t) P^{\pm}(t)\right]^l = \left[K_1^{\pm}(\beta)\right]^l, \tag{2.18}$$

so we can replace $\tilde{K}_l^{\pm}(\beta)$ by $[\tilde{K}_1^{\pm}(\beta)]^l$ in the expression for ΔS_n^{NR} to obtain that for the resetting model, ΔS_n^R. We then have

$$\Delta S_3^R = \lambda_p' A_+ \left[2 - \tilde{K}_1^-(\beta)\right] \tilde{K}_1^+(\beta) - \lambda_\pi' A_- \left[2 - \tilde{K}_1^+(\beta)\right] \tilde{K}_1^-(\beta), \tag{2.19}$$

and

$$\Delta S_4^R = \lambda_p' A_+ \left[3 - 2\tilde{K}_1^-(\beta) + \tilde{K}_1^+(\beta)\tilde{K}_1^-(\beta)\right] \tilde{K}_1^+(\beta) - \lambda_\pi' A_- \left[3 - 2\tilde{K}_1^+(\beta) + \tilde{K}_1^+(\beta)\tilde{K}_1^-(\beta)\right] \tilde{K}_1^-(\beta). \tag{2.20}$$

Clearly, as n increases, ΔS_n increases. In general, in both forms of the model, more spikes should induce greater synaptic change (ignoring issues such as saturation of synaptic strengths). It is therefore convenient to consider $\Delta \hat{S}_n = \Delta S_n / n$, and, in particular, to determine the asymptotic form as $n \to \infty$. Although the expected divergence of ΔS_n as $n \to \infty$ may seem problematic, we will see that ΔS_n possesses a fixed-point structure such that the afferents evolve to these fixed points and remain there. Hence, the divergence of ΔS_n as $n \to \infty$ is not, in practice, observed, and scaling ΔS_n by $1/n$ to remove the divergent behavior amounts simply to rescaling time or the overall learning rate.

A general analysis of n-spike trains is presented elsewhere (Appleby & Elliott, 2006). We show that by viewing the behavior of the switch as a sequence of transitions back to *OFF*, the underlying Markovian nature of the switch may be exploited to produce a generating function for any n-spike rule. The switch is Markovian in the sense that after any spike sequence that returns the switch to *OFF*, the future behavior of the switch is independent of how the switch arrived at that point. We may therefore divide any spike sequence into a chain of successive Markovian steps, each ending in a transition to *OFF*, that may be computed explicitly. Each Markovian step has an associated expected change in synaptic strength, $\mathcal{T}_1$, calculated by

conditioning on the state of the synapse as each spike arrives, as well as an expected number of spikes that occurs during the step, $\mathcal{N}_1$. For a long sequence of n spikes, we therefore expect on average $n/\mathcal{N}_1$ Markovian steps, each inducing on average a change in synaptic strength of $\mathcal{T}_1$. Hence, the total change, ΔS_n, for large enough n, goes like $n\mathcal{T}_1/\mathcal{N}_1$, and so we have, asymptotically,

$$\Delta \hat{S}_n = \frac{1}{n}\Delta S_n \sim \frac{\mathcal{T}_1}{\mathcal{N}_1}. \tag{2.21}$$

By direct computation of $\mathcal{T}_1$ and $\mathcal{N}_1$ (Appleby & Elliott, 2006), we find that the nonresetting form of the ∞-spike rule is given by

$$\Delta \hat{S}_\infty^{NR} = \frac{\lambda_\pi' A_+ K_1^+(\lambda_p) - \lambda_p' A_- K_1^-(\lambda_\pi)}{1 + \frac{\lambda_p}{\lambda_\pi} K_1^-(\lambda_\pi) + \frac{\lambda_\pi}{\lambda_p} K_1^+(\lambda_p)}, \tag{2.22}$$

and the resetting form is

$$\Delta \hat{S}_\infty^{R} = \frac{\lambda_p' A_+ [1 - \tilde{K}_1^-(\beta)]\tilde{K}_1^+(\beta) - \lambda_\pi' A_- [1 - \tilde{K}_1^+(\beta)]\tilde{K}_1^-(\beta)}{1 - \tilde{K}_1^+(\beta)\tilde{K}_1^-(\beta)}. \tag{2.23}$$

Equations 2.22 and 2.23 describe the expected change in synaptic strength arising from two freely spiking pre- and postsynaptic neurons at firing rates λ_π and λ_p, respectively. Note that the $K_1^\pm$ are functions of λ_π and λ_p in the nonresetting, ∞-spike model rather than functions of β as elsewhere. There is no difficulty with the two limits $\lambda_\pi \to 0$ or $\lambda_p \to 0$ in the expression for $\Delta \hat{S}_\infty^{NR}$ since the functions $K_1^\pm(\lambda)$ go to zero at least as fast as λ.

Starting from the two-spike rule in equation 2.8, we previously derived two constraints on the parameters $A_\pm$, $n_\pm$ and $\tau_\pm$ (Appleby & Elliott, 2005). In the limit of large λ_π and λ_p in equation 2.8, we have that $\Delta S_2 \propto A_+ - A_-$. The sign of this expression indicates whether potentiation or depression of synaptic strengths is expected for high pre- and postsynaptic firing rates. Experimental work on long-term potentiation shows that high pre- and postsynaptic firing rates generally lead to potentiation (Sjöström, Turrigiano, & Nelson, 2001). This requires that $\Delta S_2 > 0$ for large λ_π and λ_p, so that $A_+ > A_-$. Second, a necessary condition for competitive dynamics is that a depressing regime exists, as otherwise synapses can never weaken on average. Because $\Delta S_2 = 0$ at the origin, the requirement that $\partial \Delta S_2 / \partial \lambda_\pi |_{\lambda_\pi, \lambda_p = 0} < 0$ is sufficient to ensure the existence of a depressing regime. This leads to the second constraint,

$$\gamma = \frac{A_+ n_+ \tau_+}{A_- n_- \tau_-} < 1, \tag{2.24}$$

which we interpret as depression dominating over potentiation. Although necessary, the presence of a depressing regime is not, of course, sufficient to guarantee the presence of stable, competitive dynamics. The two-spike rule is qualitatively BCM-like for $\gamma < 1$, with a depressing phase at low presynaptic firing rates followed by a transition to a potentiating phase as a threshold is passed. Although we now consider multispike rules, we retain the definition of γ given by equation 2.24. Throughout this article, we set $A_+ = 1$ and $A_- = 0.95$, in accordance with the condition that $A_+ > A_-$, and choose $n_+ = n_- \in \{1, 3\}$, as in Appleby & Elliott (2005). Setting $\tau_- = 20$ ms and choosing γ therefore determines the remaining parameter τ_+. We will consider the two-, three-, four-, and ∞-spike, $n_\pm = 1$ cases in detail and then discuss how our observations extend to other n-spike rules, and for general $n_\pm$.

3 The Two-Spike Rule and Beyond

We now examine the two-spike rule and show that it leads, without further modification, to pathological learning dynamics. We then move beyond the two-spike rule and perform an initial analysis of the multispike rules, indicating the reasons that their dynamics differ from the two-spike rule's dynamics.

3.1 Failure of the Two-Spike Rule. To study two-spike interactions, we consider a system of m afferents innervating a single target cell and explore their behaviors in the context of the averaged two-spike learning rule, ΔS_2. We label the m afferents with indices such as i and j, so that $i, j \in \{1, \ldots, m\}$. Let afferent i support l_i synapses of strength $s_{i\alpha}$, $\alpha \in \{1, \ldots, l_i\}$. Let afferent i fire at rate λ_{π_i} and the target cell fire at rate λ_p. We then define

$$\frac{ds_{i\alpha}}{dt} = \Delta S_2(\lambda_{\pi_i}, \lambda_p), \tag{3.1}$$

where we have noted the explicit dependence of ΔS_2 on the pre- and postsynaptic firing rates. Since ΔS_2 is independent of the synapse label α, all $s_{i\alpha}$ experience the same change at each time step: all of afferent i's synapses experience identical pre- and postsynaptic firing rates. Hence, we may consider the evolution of either the total synaptic strength supported by afferent i, $s_i^T \equiv \sum_\alpha s_{i\alpha}$, or the average synaptic strength, $s_i^A \equiv \frac{1}{l_i}\sum_\alpha s_{i\alpha}$, evolving according to $ds_i^T/dt = l_i\, \Delta S_2(\lambda_{\pi_i}, \lambda_p)$ or $ds_i^A/dt = \Delta S_2(\lambda_{\pi_i}, \lambda_p)$, respectively. We set the postsynaptic firing rate λ_p to be the standard linear sum of the presynaptic firing rates weighted by the synaptic strengths. Thus, if we consider the total strengths s_i^T, we set $\lambda_p = \sum_i s_i^T \lambda_{\pi_i}$, while if we consider the average strengths s_i^A, we set $\lambda_p = \sum_i l_i s_i^A \lambda_{\pi_i}$. Since we study competitive dynamics, we do not consider scenarios in which some afferents at least initially enjoy an advantage over other afferents, and in particular, we do not consider scenarios in which the numbers of synapses supported by a

group of afferents differ significantly. Hence, for convenience we take l_i to be independent of i: all afferents support the same number of synapses. We may then dispense with the factors of l_i since they can be absorbed into a redefinition of time (for the total strength) or a redefinition of afferent rates, which is equivalent to a redefinition of time (for the average strength). We therefore consider one common quantity s_i, which can be thought of as either total or average synaptic strength, evolving according to the equation

$$\frac{ds_i}{dt} = \Delta S_2(\lambda_{\pi_i}, \lambda_p), \tag{3.2}$$

for which we set $\lambda_p = \sum_i s_i \lambda_{\pi_i}$. As we consider only excitatory synapses, we insist that $s_i \geq 0$. Hence, if equation 3.2 tries to drive s_i negative, we truncate the evolution and set s_i to zero. The synapse is not frozen there: it can regrow to nonzero strength.

We may obtain a qualitative understanding of the dynamics induced by equation 3.2. Consider a scenario in which at least one afferent, say afferent i, has a large synaptic strength, s_i. The postsynaptic firing rate λ_p will thus typically be large, and so the variables $\beta_i = \lambda_{\pi_i} + \lambda_p$ will typically be large. In this limit, $K_1^{\pm}(\beta_i) \approx 1$, so the two-spike learning rule reduces to $ds_j/dt \propto A_+ - A_- > 0$ for any afferent j. Thus, we see that once one afferent becomes strong, it induces all afferents to become strong, and this induces a positive feedback mechanism in which all afferents reinforce each other. So unless the first afferent that becomes strong is silenced with low λ_{π_i} for a sustained period, so that λ_p does not become large, all afferents' strengths will escape to infinity.

Consider now a scenario in which all synaptic strengths s_i are small. Here we may write $\beta_i \approx \lambda_{\pi_i}$ since λ_p is small, and the two-spike rule becomes $ds_j/dt \propto A_+ K_1^+(\lambda_{\pi_j}) - A_- K_1^-(\lambda_{\pi_j})$. Writing this out according to the definitions of $K_1^{\pm}$, we have two terms: a negative term that goes like $\gamma - 1$ and a positive term that grows with λ_{π_j}. However, only for presynaptic firing rates in excess of at least 100 Hz does ds_j/dt become overall positive. Thus, if the s_i are all small, then $ds_i/dt < 0$ for all but high, sustained firing rates, and so the s_i become even smaller. These weak afferents become trapped in the depressing phase and fall to zero strength. These simple limiting scenario arguments therefore suggest that the phase space for the two-spike rule is partitioned into two regimes, with all afferents either pulled toward zero on average or pushed toward infinity on average. Although these arguments are not rigorous proofs, their conclusion is confirmed by a full fixed-point analysis of the two-spike rule.

We therefore see that despite having a BCM-like form, the learning behavior of the rate-based, two-spike rule is in fact pathological and always leads to the afferents' either all dying or all escaping to infinity. To demonstrate that the spike-based system also exhibits the behavior characteristic of the rate-based system, in Figure 2 we show a spike-based simulation

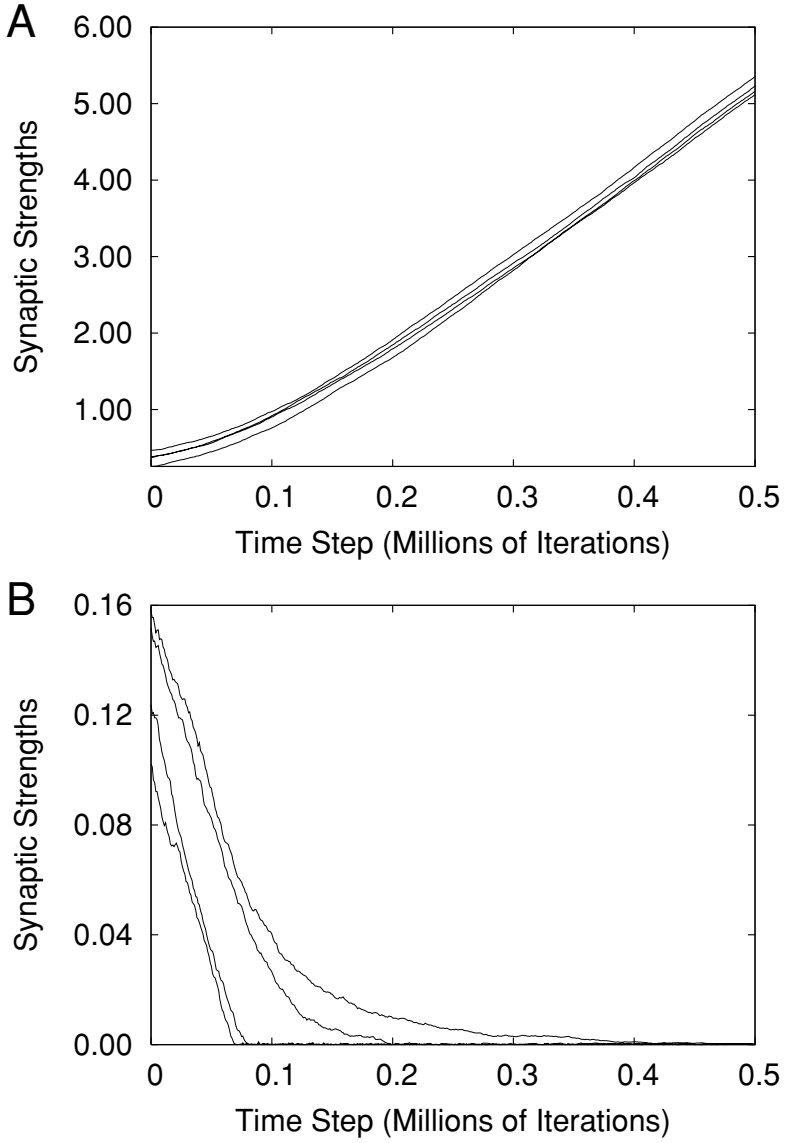

Figure 2: A spike-based simulation of the two-spike rule for four afferents innervating one target cell. The dynamics are partitioned into two distinct regimes, determined by the postsynaptic firing rate. (A) When initial synaptic strengths are large, the afferents drive the postsynaptic cell to a high firing rate, and runaway learning ensues. (B) When initial synaptic strengths are small, the postsynaptic firing rate is low, and all four afferents fall to zero strength.

of four afferents innervating a single target cell. Spike trains are truncated at two spikes: there is no interaction between successive pairs of spikes. We see that the spike-based simulation exhibits the same two regimes discussed above for the rate-based system. This behavior is, as argued above, independent of the number of afferents because the governing factor in this behavior is the postsynaptic firing rate, λ_p, which is common to all afferents synapsing on the target cell. Once λ_p begins to move toward or away from zero, uncontrolled learning ensues and the afferents either all die or all escape to infinity. The instability inherent in the two-spike learning rule, whether rate or spike based, thus shows that it is unable to support the stable, competitive dynamics characteristic of, for example, ODC formation.

This situation appears rather unpromising. For both experimental and theoretical reasons, we require that a BCM-like form emerges from the two-spike learning rule on average, yet this requirement leads directly to these pathological learning behaviors. Without further modification designed specifically to prevent runaway learning, such as hard upper bounds on synaptic strengths, the rule will always lead to uncontrolled learning. One possible remedy is to allow the threshold between the potentiating and depressing regimes, which is a function of various easily modifiable parameters, to depend on the recent time average of postsynaptic firing in a manner similar to the BCM rule (Bienenstock et al., 1982). In effect, this couples potentiation and depression together, in the sense that the dependence of the plasticity threshold on the recent time average of postsynaptic firing allows the history of potentiation and depression to influence later plasticity events. The result of this coupling in the BCM rule is to stabilize the learning behavior and prevent runaway learning.

We see a similar result when we modify the two-spike rule to incorporate a sliding threshold. In the BCM model, the sliding threshold Θ_M is explicitly set as a function of the recent time-averaged postsynaptic firing rate, $\bar{\lambda}_p$. In our two-spike rule, the threshold is dynamically determined by the solution of

$$A_+ K_1^+(\beta) = A_- K_1^-(\beta). \tag{3.3}$$

With the six parameters $A_\pm$, $n_\pm$, and $\tau_\pm$ held fixed, this gives a value for β. However, for a given value of β (and fixed parameters $A_\pm$, $n_\pm$, and τ_-), we can instead regard this equation as determining a value of τ_+. We denote this value of τ_+ by $\tau_M = f(\beta)$, where f is the function that gives the solution of equation 3.3. Since it is unrealistic for the value of τ_+ to depend on the instantaneous value of β, we instead make it depend on the recent time average, similar to the BCM rule, so that $\tau_M = f(\bar{\beta})$. Thus, for a given value of $\bar{\beta}$, our preferred value of τ_+ should be set so that $\tau_+ = \tau_M$. Such a value would place the threshold at exactly the right location, putting some afferents in the depressing region and the others in the potentiating region.

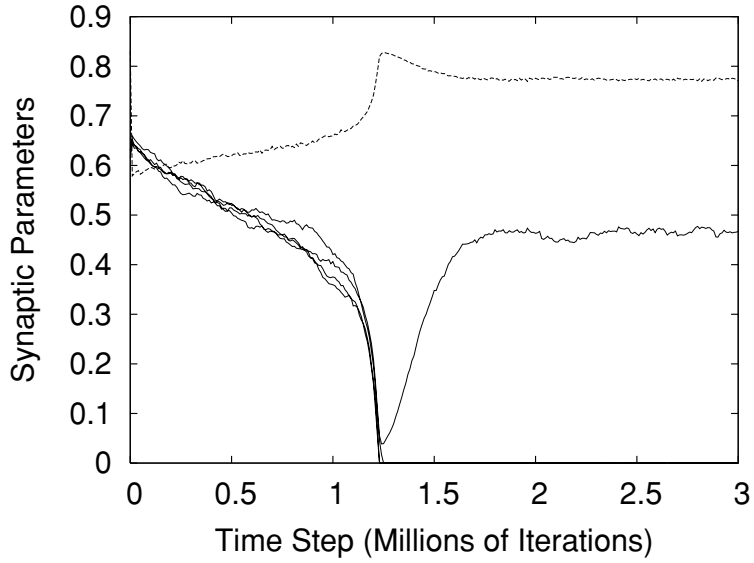

Figure 3: The evolution of four afferents innervating a single target cell under the modified two-spike, rate-based rule, which allows the threshold between depression and potentiation to slide as a function of the time-averaged postsynaptic firing rate. The four solid lines show the strengths of the four afferents, and the dashed line shows the value of γ. The introduction of the sliding threshold has stabilized the dynamics and generated competition.

Thus, if we make τ_+ dynamically evolve toward $\tau_M(\bar{\beta})$, we will achieve a sliding threshold sensitive to the recent time average of postsynaptic firing, $\bar{\lambda}_p$, through $\bar{\beta} = \bar{\lambda}_\pi + \bar{\lambda}_p$. Setting

$$\frac{d\tau_+}{dt} = \epsilon \left[\tau_M(\bar{\beta}) - \tau_+\right] \tag{3.4}$$

represents one simple way in which to realize such sliding, where ϵ is some small, inverse time constant for the relaxation of τ_+ to the (changing) value of $\tau_M(\bar{\beta})$. Implementing this sliding threshold in the two-spike rule, we replicate the dynamics of the BCM rule. Synaptic strengths are therefore stabilized, and uncontrolled learning prevented. The behavior of a set of four afferents operating under this modified, rate-based two-spike rule is shown in Figure 3. We have chosen initial conditions so that all synapses would grow without bound under the unmodified two-spike rule. For the modified rule, we see that γ therefore decreases very rapidly (within the first few thousand time steps), moving all synapses

into the depressing regime. As they depress, γ slowly increases. Eventually three of the synapses hit zero strength. The surviving synapse initially continues in the downward direction, but γ increases rapidly, moving the synapse into the potentiating regime. The synapse grows and stabilizes, and γ remains approximately constant. We therefore see, as expected, that the introduction of a sliding threshold stabilizes the learning behavior.

3.2 Beyond Two Spikes. Although coupling potentiation and depression in the manner described above for the modified two-spike rule stabilizes the learning dynamics exhibited by it, doing so forces us to make assumptions concerning the exact dependence of our model's parameters on the recent firing history. To avoid such somewhat ad hoc complications, we seek to determine whether higher-order spike interactions can instead exhibit the required stable, competitive dynamics. One reason that higher-order spike interactions might achieve this is that they automatically provide the coupling between potentiation and depression that leads to stable behavior under the modified two-spike rule. This coupling takes place in the sense that once the synapse enters the *POT* (*DEP*) state, it cannot subsequently enter the *DEP* (*POT*) state without having first returned to the *OFF* state. That is, once a synapse has entered the potentiating mode, that synapse is prevented from entering the depressing mode without having first been deactivated to the resting state. For this coupling to be expressed, however, we require a minimum of three spikes. Consider, for example, the two-spike train πp. Once the first presynaptic spike has elevated the synapse into the *POT* state, the paucity of spikes prevents the synapse from subsequently visiting the opposite half of the switch and undergoing depression. There is therefore never any coupling under two-spike trains. This is not the case for higher-order spike trains. Consider, for example, the spike train $\pi p \pi$. As before, the first presynaptic spike elevates the synapse into the *POT* state. If a transition back to *OFF* does not occur before the p spike arrives, then a potentiation event occurs as usual, and the last π event is of no importance. If, however, a transition back to *OFF* occurs before the second p spike arrives, then that p spike will cause the synapse to move to the *DEP* state, enabling the synapse possibly to undergo a depression event (assuming the final π spike arrives in good time). A similar argument applies for any number of spikes greater than two. Hence, if the switch is unified, all multispike rules couple potentiation and depression together, but the two-spike rule does not. Given that the coupling of potentiation and depression endows the BCM rule with a structure that supports stable, competitive dynamics, it is natural to ask whether the presence of this coupling in the multispike rules alters the learning dynamics compared to the two-spike rule.

In studying the dynamics of the multispike interactions, we again consider a system of m afferents innervating a single target cell and explore

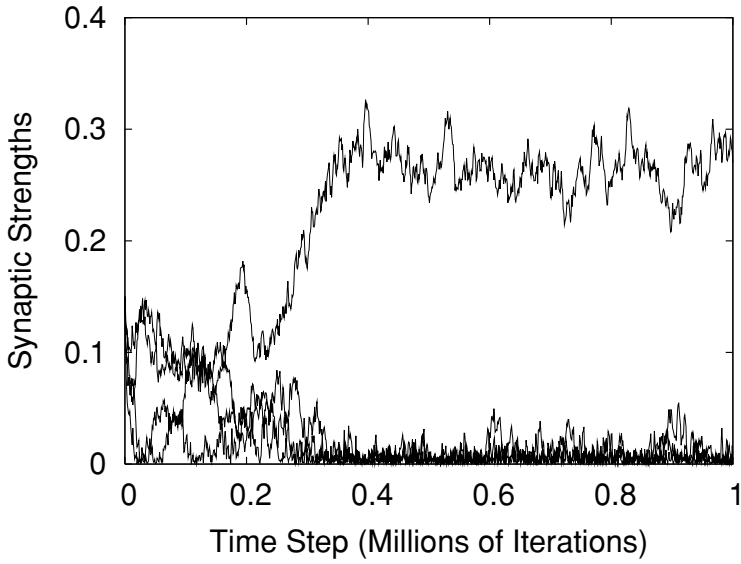

Figure 4: A spike-based simulation of the nonresetting three-spike rule for four afferents innervating one target cell. The rule is competitive, with stable segregation of the afferents being observed. This is in contrast to the two-spike rule, which requires explicit modification to achieve stable, competitive dynamics.

their behavior in the context of the averaged, n-spike learning rules. The same arguments that led to equation 3.2 now lead to

$$\frac{ds_i}{dt} = \Delta S_n(\lambda_{\pi_i}, \lambda_p) \tag{3.5}$$

as the synaptic strength modification equation corresponding to the n-spike rule. We continue to truncate s_i at zero if it is driven negative.

A numerical exploration of the three-spike, rate-based rule (either resetting or nonresetting) shows that, indeed, its learning dynamics differ significantly from those of the two-spike rule. Under the three-spike rule, the uncontrolled learning behavior seen for the (unmodified) two-spike rule is absent for a broad range of parameters. Afferents compete for control of the target cell, with stable segregation robustly occurring. Figure 4 confirms that these observations are also true for the spike-based version of the three-spike rule. Stable, segregated fixed points exist under the three-spike rule, leading to the same competitive dynamics exhibited by the BCM rule. Thus, the ability of the three-spike rule to probe the coupling of potentiation and depression under a unified switch, however modest,

dramatically alters the dynamical landscape compared to that of the two-spike rule, which cannot probe coupling. We observe a similar result for any multispike rule. Attempts to modify the two-spike rule, such as introducing a sliding threshold, while adequate in terms of stabilizing the learning dynamics, are therefore unnecessary: all that is required is to extend our consideration of spike interactions to three or more spikes with no ad hoc modifications of the learning rules.

The success of the multispike rules depends critically on the unification of the switch mechanism and thus the coupling of potentiation and depression. We can see this explicitly by examining the multispike learning rules for two ununified switches, so that we consider the potentiating lobe in Figure 1A and the depressing lobe in Figure 1B separately. We can reduce the unified switch to two separate potentiating and depressing switches by setting $\tau_- = 0$ and $\tau_+ = 0$, respectively. Setting $\tau_- = 0$, for example, means that the depressing lobe of the unified switch is effectively unavailable, because any transition to the *DEP* state results in an instantaneous stochastic decay back to *OFF* with no depression. Similarly with $\tau_+ = 0$, the potentiating lobe is unavailable. Adding the two rules derived by separately setting $\tau_+ = 0$ and $\tau_- = 0$ gives the plasticity rule corresponding to the operation of two ununified, separate switches. Consider, for example, the resetting ∞-spike rule given in equation 2.23. Setting $\tau_- = 0$ gives

$$\Delta \hat{S}_\infty^{POT} = \lambda'_p A_+ \tilde{K}_1^+(\beta), \tag{3.6}$$

and setting $\tau_+ = 0$ gives

$$\Delta \hat{S}_\infty^{DEP} = -\lambda'_\pi A_- \tilde{K}_1^-(\beta), \tag{3.7}$$

so that the overall rule is

$$\Delta \hat{S}_\infty^{POT} + \Delta \hat{S}_\infty^{DEP} = \lambda'_p A_+ \tilde{K}_1^+(\beta) - \lambda'_\pi A_- \tilde{K}_1^-(\beta) \equiv \Delta S_2. \tag{3.8}$$

The resetting ∞-spike rule therefore reduces to the two-spike rule when the switch is split into two halves. Repeating this manipulation for any resetting multispike rule produces the same result, so that $\Delta S_n^{POT} + \Delta S_n^{DEP} = \Delta S_2 \, \forall n$. Thus, the unification of the separate switches, which was proposed initially on the grounds of simplicity, has unexpected consequences for the dynamics of the model. Examining the nonresetting ∞-spike rule under an identical manipulation, we find that the potentiating half (when $\tau_- = 0$) gives

$$\Delta \hat{S}_\infty^{POT} = \frac{\lambda'_\pi A_+ K_1^+(\lambda_p)}{1 + \frac{\lambda_\pi}{\lambda_p} K_1^+(\lambda_p)}, \tag{3.9}$$

and the depressing half (when $\tau_+ = 0$) gives

$$\Delta \hat{S}_\infty^{DEP} = -\frac{\lambda'_p A_- K_1^-(\lambda_\pi)}{1 + \frac{\lambda_p}{\lambda_\pi} K_1^-(\lambda_\pi)}, \tag{3.10}$$

so that the overall rule is

$$\Delta \hat{S}_\infty^{POT} + \Delta \hat{S}_\infty^{DEP} = \frac{\lambda'_\pi A_+ K_1^+(\lambda_p)}{1 + \frac{\lambda_\pi}{\lambda_p} K_1^+(\lambda_p)} - \frac{\lambda'_p A_- K_1^-(\lambda_\pi)}{1 + \frac{\lambda_p}{\lambda_\pi} K_1^-(\lambda_\pi)}. \tag{3.11}$$

Although this nonresetting rule has not reduced to the two-spike rule, a similar analysis of the cases of large and small s_i (or large and small λ_p) as performed above for the two-spike rule reveals identical conclusions, so that afferents either all escape to infinity or all die at zero. We therefore conclude that the presence of higher-order spike interactions under a unified three-state synaptic switch differentiates the two-spike and multispike rules by allowing a probing of the coupling between potentiation and depression in the unified model. These higher-order interactions are responsible for giving rise to the stable, competitive dynamics that we observe for the multispike rules.

Examining the large β limit of the (unified) multispike rules reveals that their asymptotic behavior differs significantly from that of the two-spike rule. For large β, $K_l^{\pm}(\beta) \to 1$, so $\tilde{K}_l^+(\beta) \to {\lambda'_\pi}^l$ and $\tilde{K}_l^-(\beta) \to {\lambda'_p}^l$. Because $\lambda'_\pi + \lambda'_p = 1$, we introduce a new variable, $x \in [-1, +1]$, such that

$$\lambda'_\pi = \frac{1}{2}(1 - x), \tag{3.12}$$

$$\lambda'_p = \frac{1}{2}(1 + x). \tag{3.13}$$

It is easy to see that x is the tangent of the angle between the vector $(\lambda_\pi, \lambda_p)^T$, the superscript T denoting the transpose, and the line $\lambda_\pi = \lambda_p$. If θ is the standard angle in a polar coordinate system ($\lambda_\pi = r\cos\theta, \lambda_p = r\sin\theta$), then x is just

$$x = \tan(\theta - \pi/4). \tag{3.14}$$

In this large β limit, we rewrite the two- and multispike rules in terms of the variable x and find that

$$\Delta S_2 \to \frac{1}{4}(1 - x^2)(A_+ - A_-), \tag{3.15}$$

$$\Delta S_3 \to \frac{1}{8}(1-x^2)\left[A_+(3-x) - A_-(3+x)\right], \tag{3.16}$$

$$\Delta S_4 \to \frac{1}{16}(1-x^2)\left[A_+(9-4x-x^2) - A_-(9+4x-x^2)\right], \tag{3.17}$$

$$\Delta \hat{S}_\infty \to \frac{1}{2}\left(\frac{1-x^2}{3+x^2}\right)\left[A_+(1-x) - A_-(1+x)\right]. \tag{3.18}$$

Because the $\beta \to \infty$ limits of the resetting and nonresetting models are identical, the results above are independent of the form of the model, although the limit of the ∞-spike resetting rule is much easier to extract. We see that while the two-spike rule, ΔS_2, is symmetrical about $x = 0$, and thus symmetrical about the line $\lambda_\pi = \lambda_p$, the other rules exhibit an asymmetry about $x = 0$ due to the presence of odd powers of x. This property is true for all the multispike rules, not just ΔS_3, ΔS_4, and $\Delta \hat{S}_\infty$ given above. For $x^2 < 1$, the right-hand side of equation 3.14 is strictly positive, because $A_+ > A_-$. Hence, the two-spike rule always potentiates in the large β limit, as we saw earlier. This behavior is not the case for the multispike rules. Consider the cases $x \approx +1$ and $x \approx -1$ in the above. Then we find

$$\Delta S_3 \propto \begin{cases} A_+ - 2A_- & \text{for } x \approx +1 \\ 2A_+ - A_- > 0 & \text{for } x \approx -1 \end{cases} \tag{3.19}$$

$$\Delta S_4 \propto \begin{cases} A_+ - 3A_- & \text{for } x \approx +1 \\ 3A_+ - A_- > 0 & \text{for } x \approx -1 \end{cases} \tag{3.20}$$

$$\Delta \hat{S}_\infty \propto \begin{cases} -A_- < 0 & \text{for } x \approx +1 \\ +A_+ > 0 & \text{for } x \approx -1 \end{cases} \tag{3.21}$$

Indeed, by using the general form of the $(n+1)$-spike rule (Appleby & Elliott, 2006), we find

$$\Delta S_{n+1} \propto \begin{cases} A_+ - nA_- & \text{for } x \approx +1 \\ nA_+ - A_- > 0 & \text{for } x \approx -1 \end{cases}. \tag{3.22}$$

The multispike rules therefore always potentiate in the $x \approx -1$ or large λ_π, small λ_p direction. For the three-spike rule, if $A_+ < 2A_-$, then it, and all higher-order rules, depress in the $x \approx +1$ or small λ_π, large λ_p direction. If, however, $A_+ > 2A_-$, then the three-spike rule potentiates in all directions and will exhibit runaway learning just like the two-spike rule. Although it may be the case that $A_+ > 2A_-$, it may not be the case that $A_+ > 3A_-$. Here, although the three-spike rule fails, the four-spike rule will depress in the $x \approx +1$ direction. Indeed, by looking at the behavior of the general $(n+1)$-

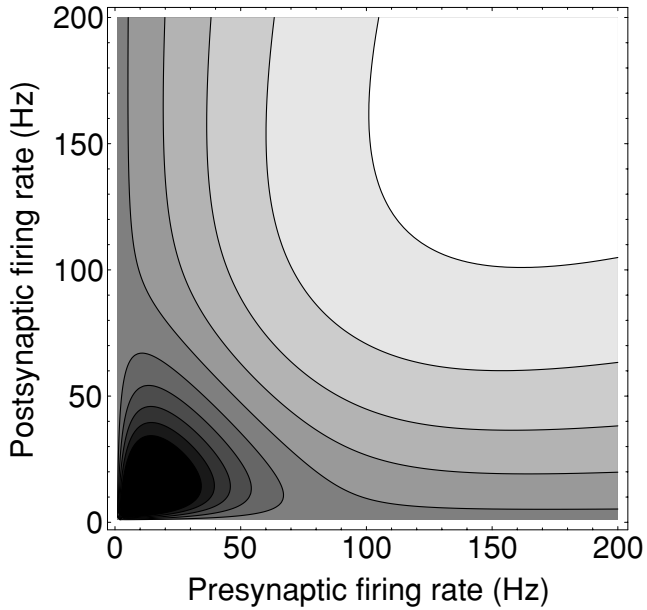

Figure 5: A contour plot of ΔS_2 in the λ_π–λ_p plane. Black areas represent minimum values and white areas maximum values, and nine shades of gray interpolate between these extremes. The minimum value on this partial plane is -0.0244, and the maximum value is $+0.0118$.

spike rule in equation 3.22, we see that provided $A_- \neq 0$, there always exists a value of n above which the multispike rules will start to depress in the $x \approx +1$ direction. In particular, for $n > A_+/A_-$, the $(n+1)$-spike rule will always potentiate in the $x \approx -1$ direction and depress in the $x \approx +1$ direction. With such "mixed" dynamics at large β, large λ_p will not induce runaway learning. Since $A_+/A_- > 1$, the two-spike rule can never achieve this. Hence, the two-spike rule is irredeemably pathological in its learning behavior due to a symmetry that is absent in all the multispike learning rules, and although the multispike rules can fail in the same way as the two-spike rule, this is parameter dependent (unlike the two-spike rule), and we are guaranteed (for $A_- \neq 0$) that there exists an n above which the n-spike rules will not fail.

Plotting the ΔS_n in the λ_π–λ_p plane illustrates these results. For the two-spike rule (see Figure 5), we see a depressing well around the origin and a potentiating regime away from it. The two-spike surface is symmetrical about the line $\lambda_\pi = \lambda_p$. Depression therefore always occurs at low β, and potentiation always occurs at high β. For the nonresetting forms of the three- (see Figure 6), four- (see Figure 7), and ∞-spike (see Figure 8) rules,

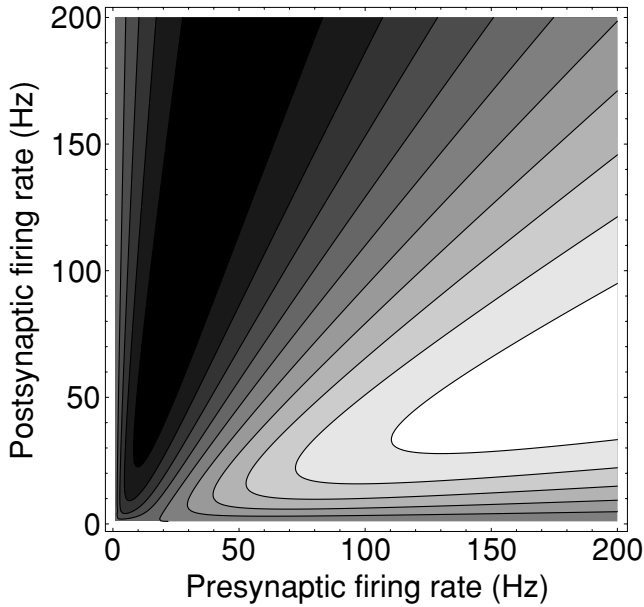

Figure 6: A contour plot of ΔS_3^{NR} for the nonresetting model in the λ_π–λ_p plane. The minimum value on this partial plane is -0.0826, and the maximum value is $+0.1007$.

the symmetry about the line $\lambda_\pi = \lambda_p$ is absent, and it is possible to induce either potentiation or depression at high β depending on the value of θ (or x), for the parameter choice used here. Thus, the values of λ_π and λ_p together determine whether potentiation or depression occurs.

4 Fixed-Point Analysis of a Rate-Based Rule

Having derived various multispike, rate-based rules from our switch model and performed an initial study of the differences between the two- and multispike rules, we may proceed to develop a deeper analytical understanding of the learning dynamics exhibited by the multispike rules. In particular, we continue to study equation 3.5 by performing a fixed-point analysis of the ∞-spike learning rule.

We assume that all m afferents have the same mean firing rate, $\mu > 0$, and variance, $\sigma^2 > 0$. The firing rate of each afferent therefore fluctuates about a common mean. The fluctuations distinguish the afferents (unless they are perfectly correlated), while preventing any afferent from enjoying an overall advantage. We set $\lambda_{\pi_i} = \mu(1 + \alpha_i)$, where α_i is some small perturbation about the mean, $|\alpha_i| \ll 1$, and take the mean of α_i to be zero, $\langle \alpha_i \rangle = 0$, so

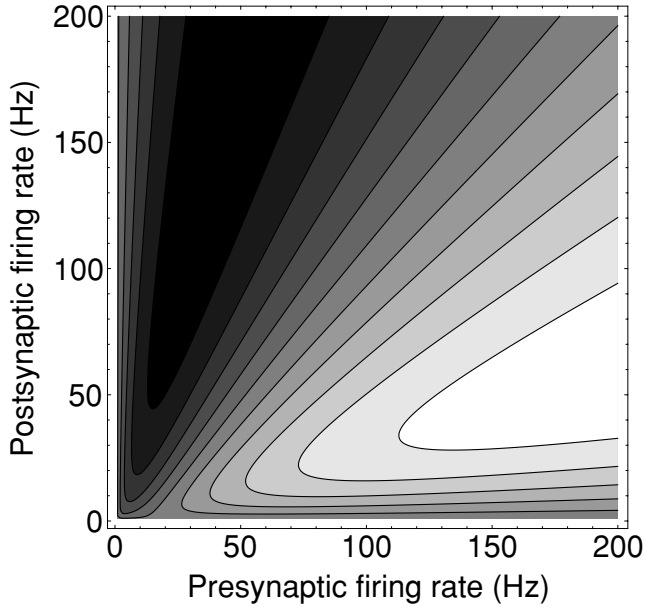

Figure 7: A contour plot of ΔS_4^{NR} for the nonresetting model in the λ_π–λ_p plane. The minimum value on this partial plane is -0.1686, and the maximum value is $+0.1926$.

that $\langle \lambda_{\pi_i} \rangle = \mu$, as required. As we will average over the ensemble of activity patterns, we must obtain an expression for $\langle \alpha_i \, \alpha_j \rangle$. The variance in the activity of afferent i is

$$\sigma^2 = \left\langle \lambda_{\pi_i}^2 \right\rangle - \langle \lambda_{\pi_i} \rangle^2 = \mu^2 \left\langle \alpha_i^2 \right\rangle \tag{4.1}$$

so that $\langle \alpha_i^2 \rangle = \hat{\sigma}^2$, where $\hat{\sigma} = \sigma/\mu$, with $\hat{\sigma}^2 \ll 1$. Assuming for convenience that the afferents' activities are uncorrelated, so that their covariance $\mathrm{Cov}(\alpha_i, \alpha_j) = 0$ for $i \neq j$, we then have

$$\langle \alpha_i \, \alpha_j \rangle = \hat{\sigma}^2 \delta_{ij}, \tag{4.2}$$

where δ_{ij} is the Kronecker delta, equal to one if $i = j$ and zero otherwise. Defining the vectors $\underline{s} = (s_1, \ldots, s_m)^T$ and $\underline{\alpha} = (\alpha_1, \ldots, \alpha_m)^T$, we then have

$$\langle \underline{\alpha} \cdot \underline{s} \rangle = 0, \tag{4.3}$$

$$\langle \alpha_i (\underline{\alpha} \cdot \underline{s}) \rangle = \hat{\sigma}^2 s_i, \tag{4.4}$$

$$\langle (\underline{\alpha} \cdot \underline{s})^2 \rangle = \hat{\sigma}^2 |\underline{s}|^2. \tag{4.5}$$

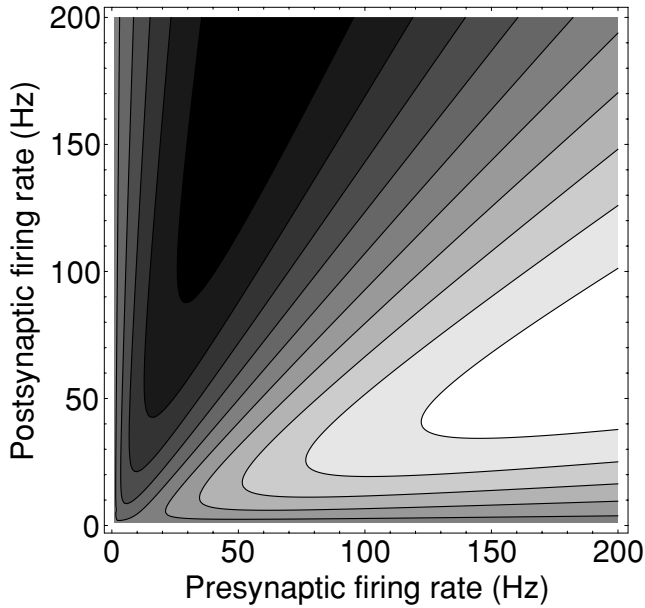

Figure 8: A contour plot of $\Delta\hat{S}_{\infty}^{NR}$ for the nonresetting model in the λ_{π}–λ_{p} plane. The minimum value on this partial plane is -0.1010, and the maximum value is $+0.1036$.

As the perturbations are small, $\hat{\sigma}^2 \ll 1$, we may expand any n-spike rule in α_i and then average over the ensemble of afferent activity patterns using the three equations above. This expansion must be to second order in α_i, as the mean of α_i is zero. We thus arrive at a set of equations describing the evolution of afferents governed by the n-spike switch rule when the activities of the afferents fluctuate about some common mean firing rate. We may then extract the fixed-point structure that characterizes the dynamics of this system. Because $s_i \geq 0$, any fixed points must lie in the nonnegative hyperquadrant of the vector space defined by $\underline{s}$ in order to be accessible to the afferents.

We now explore the fixed-point structure of a simplified form of the ∞-spike rule, the simplification merely allowing a less messy analytical characterization of the locations and stabilities of the various fixed points. We then examine the full, unsimplified rule. Although analytical results can still be obtained for the full rule, they are messy, cumbersome, and rather opaque, so we do not reproduce them here. Nevertheless, both the simplified and full forms of the model exhibit qualitatively identical dynamics.

4.1 Simplified ∞-Spike Rule. We consider a simplified form of the full ∞-spike nonresetting rule given by equation 2.22 that excludes the denominator. This exclusion simplifies the expansion, making the resulting expressions more transparent. The price for this transparency is that the location and stability of any fixed point in this simplified model will be slightly different from that of the full model. To zeroth order in $\hat{\sigma}^2$, the simplified model and full model are, however, identical in their fixed-point structure. For reasons of analytical tractability, we set $n_\pm = 1$. The simplified ∞-spike nonresetting rule is then

$$\Delta \hat{S}_\infty^{Sim}(\lambda_\pi, \lambda_p) = \lambda'_\pi A_+ K_1^+(\lambda_p) - \lambda'_p A_- K_1^-(\lambda_\pi), \tag{4.6}$$

and we set

$$\frac{ds_i}{dt} = \Delta \hat{S}_\infty^{Sim}(\lambda_{\pi_i}, \lambda_p). \tag{4.7}$$

We denote the sum of the synaptic weights as $s_+ = \sum_i s_i$, and we define $t_+ = 1 + s_+$. Expanding the simplified rule to second order in α_i and averaging over the ensemble of afferent activity patterns as set out above yields the averaged form of the rule. Dropping the angle brackets around the s_i for notational convenience, after lengthy algebra we obtain

$$\frac{ds_i}{dt} = A_- \tau_- \mu \frac{1}{t_+} \left(s_+ N_0 + \hat{\sigma}^2 N_1^i \right), \tag{4.8}$$

where

$$N_0 = \frac{\gamma}{1 + \mu \tau_+ s_+} - \frac{1}{1 + \mu \tau_-}, \tag{4.9}$$

and

$$N_1^i = \frac{\gamma}{1 + \mu \tau_+ s_+} X_i - \frac{1}{1 + \mu \tau_-} Y_i, \tag{4.10}$$

where

$$\begin{aligned} X_i = {} & \frac{1}{1 + \mu \tau_+ s_+} \left(s_i - \frac{s_i + |\underline{s}|^2}{t_+} \right) + s_+ \frac{1 + 2s_i + |\underline{s}|^2}{t_+^2} \\ & - s_+ \frac{1 + s_i}{t_+} - \frac{\mu \tau_+ |\underline{s}|^2}{(1 + \mu \tau_+ s_+)^2}, \end{aligned} \tag{4.11}$$

and

$$Y_i = \frac{1}{1+\mu\tau_-}\left(s_i - s_+\frac{1+s_i}{t_+}\right) + s_+\frac{1+2s_i+|\underline{s}|^2}{t_+^2} - \frac{s_i+|\underline{s}|^2}{t_+} - \frac{\mu\tau_- s_+}{(1+\mu\tau_-)^2}. \tag{4.12}$$

At a fixed point, we require that $d\underline{s}/dt = \underline{0}$. Solving this equation exactly for the location of all the fixed points is usually difficult, if not impossible, so we proceed by finding an approximation to zeroth order in $\hat{\sigma}^2$, for which all the fixed points can be located, and then calculate first-order corrections in $\hat{\sigma}^2$. We therefore write a fixed point as

$$\underline{s}_{FP} = \underline{x} + \hat{\sigma}^2\underline{y}, \tag{4.13}$$

where $\underline{x}$ is the zeroth-order approximation to the location of the fixed point, $\underline{s}_{FP}$, and $\underline{y}$ is the first-order correction. We define $x_+ = \sum_i x_i$ and $y_+ = \sum_i y_i$, so that $s_+ = x_+ + \hat{\sigma}^2 y_+$.

4.1.1 Zeroth-Order Solutions and Behavior. To zeroth order in $\hat{\sigma}^2$, equation 4.8 becomes

$$\frac{ds_i}{dt} = A_-\tau_-\mu\frac{s_+}{t_+}\left(\frac{\gamma}{1+\mu\tau_+ s_+} - \frac{1}{1+\mu\tau_-}\right), \tag{4.14}$$

which depends only on s_+ and not the individual components s_i. By inspection, we see that there are two fixed hyperplanes. One of these corresponds to the hyperplane $x_+ = 0$. All points except $\underline{x} = \underline{0}$ in this hyperplane have at least one negative component of $\underline{x}$, and all of these points are therefore forbidden. Hence, the only fixed point on this hyperplane of fixed points accessible to the afferent weight vector is the origin. We refer to this point throughout as the *zero fixed point*. The other fixed hyperplane arises from the solution of

$$\frac{\gamma}{1+\mu\tau_+ x_+} = \frac{1}{1+\mu\tau_-}, \tag{4.15}$$

or

$$x_+ = \frac{\gamma(1+\mu\tau_-)-1}{\mu\tau_+}. \tag{4.16}$$

So that at least some points in the hyperplane $s_+ = x_+$ have all nonnegative components, we require $x_+ > 0$, or

$$\gamma > \gamma_0 \equiv \frac{1}{1 + \mu\tau_-}. \tag{4.17}$$

We refer to the hyperplane $s_+ = x_+ > 0$ as the nonzero fixed hyperplane. If $\gamma = \gamma_0$, then the nonzero hyperplane becomes coincident with the hyperplane $s_+ = 0$, and the only permitted fixed point that exists for $\gamma = \gamma_0$ is the origin, $\underline{x} = \underline{0}$.

To determine the stability of these fixed points, we examine the behavior of the system under small perturbations about them. We denote a perturbation by $\underline{\delta} = (\delta_1, \ldots, \delta_m)^T$ and write $\underline{s} = \underline{s}_{FP} + \underline{\delta}$. We define $\delta_+ = \sum_i \delta_i$.

Expanding and linearizing equation 4.14 about the zero fixed point, $\underline{s}_{FP} = \underline{0}$, we have

$$\frac{d\delta_i}{dt} = A_-\tau_-\mu(\gamma - \gamma_0)\delta_+. \tag{4.18}$$

The m eigenvalues of the associated matrix characterizing these linearized dynamics are therefore easily seen to be

$$\lambda_1 = mA_-\tau_-\mu(\gamma - \gamma_0), \tag{4.19}$$

and

$$\lambda_i = 0, \quad \forall\, i > 1. \tag{4.20}$$

The $m - 1$ repeated zero eigenvalues indicate that there is no flow along the associated eigenvectors. The flow toward or away from the origin will therefore occur only parallel to the eigenvector associated with λ_1, which is $(1, \ldots, 1)^T$. The stability of the zero fixed point is determined by the sign of λ_1. For $\gamma < \gamma_0$, the origin is stable, while for $\gamma > \gamma_0$, it is unstable. Note that when $\gamma > \gamma_0$, the nonzero fixed hyperplane $s_+ = x_+ > 0$ intersects the positive hyperquadrant, so the origin becomes unstable precisely when this hyperplane moves into the positive hyperquadrant.

Now expanding and linearizing equation 4.14 about any point in the nonzero fixed hyperplane, we obtain

$$\frac{d\delta_i}{dt} = -A_-\tau_-\tau_+\mu^2\frac{x_+}{1 + x_+}\gamma^{-1}\frac{1}{(1 + \mu\tau_-)^2}\delta_+, \tag{4.21}$$

with associated eigenvalues

$$\lambda_1 = -mA_-\tau_-\tau_+\mu^2\frac{x_+}{1+x_+}\gamma^{-1}\frac{1}{(1+\mu\tau_-)^2}, \tag{4.22}$$

and

$$\lambda_i = 0, \quad \forall\, i > 1. \tag{4.23}$$

The flow toward the nonzero fixed hyperplane is parallel to $(1,\ldots,1)^T$. As for the zero fixed point, the sign of λ_1 determines the stability of the nonzero fixed hyperplane. For $x_+ > 0$, that is, when the nonzero fixed hyperplane intersects the positive hyperquadrant, the nonzero fixed hyperplane is stable. Note that $x_+ > 0$ requires $\gamma > \gamma_0$, so when the nonzero fixed hyperplane is stable, the origin is unstable.

The zeroth-order dynamics are therefore uniquely determined by the sign of the quantity $\gamma - \gamma_0$. We briefly summarize the dynamics in each of the two possible regimes.

When $\gamma < \gamma_0$, a single fixed point $\underline{s}_{FP} = \underline{0}$ is permitted for the afferent weight vector. This point is stable, and the afferent weight vector will initially flow toward the hyperplane $s_+ = 0$ parallel to the vector $(1,\ldots,1)^T$. When it hits a hyperplane defined by $s_i = 0$, for some i, it is prevented from crossing it because s_i is always truncated at zero. The weight vector therefore remains in this $s_i = 0$ hyperplane and flows toward the origin. It may hit other $s_j = 0$, $j \neq i$, hyperplanes as it further evolves and again will be constrained to remain in them. The weight vector will therefore always arrive at the origin, regardless of the initial conditions.

When $\gamma > \gamma_0$, the origin is an unstable fixed point, and there exists a hyperplane of stable, nonzero fixed points in the positive hyperquadrant. The afferent weight vector initially flows parallel to the vector $(1,\ldots,1)^T$ toward the hyperplane $s_+ = x_+ > 0$, from either above or below it. For some initial conditions, the weight vector will directly hit the nonzero hyperplane and stop evolving. For other initial conditions, the weight vector will hit a hyperplane $s_i = 0$ for some i first, and then flow in this hyperplane until it reaches the intersection with the $s_+ = x_+ > 0$ hyperplane, and then stop evolving. Regardless of the initial conditions, the weight vector will always arrive at some point on the nonzero fixed hyperplane.

Of course, overall these dynamics are rather uninteresting, precisely because the zeroth-order solutions do not discriminate between afferents, since all afferents fire with a common rate μ. Nevertheless, the zeroth-order solutions are the foundation on which the first-order corrections are determined, and this is why we have labored the analysis of the zeroth-order case somewhat. The first-order corrections do permit a discrimination between afferents based on the fluctuations in their firing rates, and so we

expect to find a more compelling set of dynamics at first order. We now turn to this case.

4.1.2 First-Order Corrections and Behavior. We now examine the full form of equation 4.8, including the first-order corrections. For the first-order dynamics, we find that several different classes of fixed point exist.

By simple inspection of the terms in N_1^i, the first-order correction in equation 4.8, we see that the system still possesses a fixed point at $\underline{s} = \underline{0}$, and thus there are obviously no first-order corrections to its location. However, there are corrections to the eigenvalues of the stability matrix. Expanding and linearizing equation 4.8 about the zero fixed point as usual, we now find

$$\frac{d\delta_i}{dt} = A_-\tau_-\mu\left[\left(\gamma - \frac{1}{1+\mu\tau_-} - \hat{\sigma}^2\frac{\mu^2\tau_-^2}{(1+\mu\tau_-)^3}\right)\delta_+ + \hat{\sigma}^2\frac{\mu\tau_-}{(1+\mu\tau_-)^2}\delta_i\right], \tag{4.24}$$

with associated eigenvalues

$$\lambda_1 = mA_-\tau_-\mu\left[\gamma - \frac{1}{1+\mu\tau_-} - \hat{\sigma}^2\frac{\mu\tau_-}{(1+\mu\tau_-)^2}\left(\frac{\mu\tau_-}{1+\mu\tau_-} - \frac{1}{m}\right)\right] \tag{4.25}$$

and

$$\lambda_i = A_-\tau_-\mu\frac{\mu\tau_-}{(1+\mu\tau_-)^2}, \quad \forall\, i > 1. \tag{4.26}$$

The eigenvalues λ_i, $i > 1$, are always positive, so the first-order corrections have made the zero fixed point always unstable. The precise classification of the zero fixed point depends on the sign of λ_1. When $\lambda_1 < 0$, the origin is a saddle node, and when $\lambda_1 > 0$, it is a repeller. The transition of the zero fixed point from a saddle to a repeller occurs at the value of γ given by

$$\gamma = \gamma_2 \equiv \frac{1}{1+\mu\tau_-}\left[1 + \hat{\sigma}^2\frac{\mu\tau_-}{1+\mu\tau_-}\left(\frac{\mu\tau_-}{1+\mu\tau_-} - \frac{1}{m}\right)\right]. \tag{4.27}$$

Unlike the zero fixed point, the first-order corrections do affect the location of the nonzero fixed points. Indeed, the corrections destroy the entire nonzero hyperplane of fixed points, leaving one real fixed point and m quasi-fixed points, together with a set of other fixed points that are essentially uninteresting because their existence merely renders the global fixed-point structure consistent. Unless $\underline{s} = \underline{0}$, the first-order corrections in the expression for N_1^i break the symmetry between the afferents that is present at zeroth order. It is precisely this symmetry that endows the zeroth-order

system with an entire hyperplane of fixed points. At first order, this symmetry is absent, and the hyperplane collapses into a set of isolated quasi-fixed points and real fixed points. We first define what we mean by quasi-fixed points and examine their stability, then examine the other fixed points.

In models of synaptic competition that possess a fixed-point structure, there are usually fixed points in which all but one s_i are zero. Such fixed points correspond to segregated states, since at each fixed point, only one afferent innervates the target cell. With m afferents, there are m such segregated fixed points. Of course, the defining condition of a fixed point is that $d\underline{s}/dt = \underline{0}$ when the derivatives are evaluated at the fixed point. Suppose, however, that we have a point for which $s_i \neq 0$ and $s_j = 0\ \forall j \neq i$ for which $ds_i/dt = 0$ and $ds_j/dt < 0$, $\forall j \neq i$, when the derivatives are evaluated at this point. Such a point is not strictly a fixed point. However, if a model's dynamics include truncation of s_j as it tries to pass through zero into a region of negative s_j, s_j will be returned to zero. Such a point would therefore appear to be a fixed point, since the dynamics could evolve the weight vector to this point and then would remain there. We refer to such points as *quasi-fixed points*. If $ds_j/dt < 0$ for those components that are zero and if the nonzero s_i direction is stable, then we refer to the quasi-fixed point as stable; otherwise we refer to the quasi-fixed point as unstable.

Equation 4.8 possesses m such points corresponding to segregated quasi-fixed points. Consider the point $\underline{s}_{QFP} = (0, \ldots, 0, s_+, 0, \ldots, 0)^T$, where only the ith component is nonzero and where we write the usual expansion, $s_+ = x_+ + \hat{\sigma}^2 y_+$. The component s_i must therefore have zero derivative at $\underline{s}_{QFP}$, and this requirement determines the value of s_+. Solving the equation $ds_i/dt|_{\underline{s}_{QFP}} = 0$, we obtain the zeroth-order solution,

$$x_+ = \frac{\gamma(1 + \mu\tau_-) - 1}{\mu\tau_+}, \tag{4.28}$$

which is equation 4.16, reflecting the fact that all nonzero solutions at zeroth order must live in the nonzero fixed hyperplane. After a little algebra, the first-order correction is found to be

$$y_+ = \frac{1}{\mu\tau_+} \frac{1 - \gamma}{\gamma} \frac{1 - \gamma\mu\tau_-}{1 + \mu\tau_-}. \tag{4.29}$$

We require that $s_+ = x_+ + \hat{\sigma}^2 y_+ > 0$ for this point to be in the nonnegative hyperquadrant, so we will obtain a first-order correction to the bound on γ. At zeroth order, the condition that $x_+ > 0$ forces $\gamma > \gamma_0$. We now write $\gamma = \gamma_0 + \hat{\sigma}^2\hat{\gamma}$ and determine a condition on $\hat{\gamma}$ so that $s_+ > 0$. We find that

$$\hat{\gamma} > -\frac{\mu\tau_-}{(1 + \mu\tau_-)^3}, \tag{4.30}$$

and so, for $s_+ > 0$ at these m possible segregated points, we need

$$\gamma > \gamma_1 \equiv \frac{1}{1+\mu\tau_-}\left[1 - \hat{\sigma}^2 \frac{\mu\tau_-}{(1+\mu\tau_-)^2}\right]. \tag{4.31}$$

It is easy to see that $\gamma_2 \geq \gamma_1$ since $m \geq 1$. Thus, these m points become accessible to the weight vector before the origin turns from a saddle into a repeller. We now need to determine the sign of ds_j/dt, $j \neq i$, at $\underline{s}_{QFP}$ in order to determine whether these points are possibly stable. We find that for $j \neq i$,

$$\left.\frac{ds_j}{dt}\right|_{\underline{s}=\underline{s}_{QFP}} = A_-\tau_-\mu\hat{\sigma}^2 \frac{x_+}{(1+x_+)^2}\frac{1}{\gamma}\frac{1}{1+\mu\tau_-} \times \left[x_+\left(\gamma - \frac{1}{1+\mu\tau_-}\right) - \gamma\frac{\mu\tau_-}{1+\mu\tau_-}\right]. \tag{4.32}$$

This equation is purely first order in $\hat{\sigma}^2$, and thus any resulting bound on γ derived from it will be purely zeroth order in $\hat{\sigma}^2$. To calculate a higher-order correction to any resulting bound on γ, we would be compelled to extend our expansion out to order $\hat{\sigma}^4$. One consequence is that while at first order in $\hat{\sigma}^2$ the quantity x_+ is allowed to be negative, because the correction $\hat{\sigma}^2 y_+$ can pull the sum $s_+ = x_+ + \hat{\sigma}^2 y_+$ overall positive, nevertheless, in determining the sign of the right-hand side of equation 4.31, we must take x_+ strictly positive, because the resulting bound on γ will be only zeroth order, for which we are allowed to have only $x_+ > 0$. Hence, the sign of the right-hand side of equation 4.32 is determined by the terms in square brackets. For negative derivatives, we require

$$x_+\left(\gamma - \frac{1}{1+\mu\tau_-}\right) - \gamma\frac{\mu\tau_-}{1+\mu\tau_-} < 0. \tag{4.33}$$

Replacing x_+ by its expression in equation 4.28 and writing $\tau_+ = \gamma\tau_- A_-/A_+$ since we regard τ_+ as a function of γ with τ_- and $A_\pm$ fixed, we then obtain the quadratic equation

$$\left[1 - \frac{A_-}{A_+}(\mu\tau_-\gamma_0)^2\right]\gamma^2 - 2\gamma_0\gamma + \gamma_0^2 < 0. \tag{4.34}$$

Solving this equation for γ gives bounds on γ that, after a little algebra, are

$$\frac{A_+(1+\mu\tau_-) - \mu\tau_-\sqrt{A_+A_-}}{A_+(1+2\mu\tau_-) + (A_+ - A_-)\mu^2\tau_-^2} < \gamma < \frac{A_+(1+\mu\tau_-) + \mu\tau_-\sqrt{A_+A_-}}{A_+(1+2\mu\tau_-) + (A_+ - A_-)\mu^2\tau_-^2}. \tag{4.35}$$

Notice that $\gamma = \gamma_0$ satisfies equation 4.34, so the relevant bound on γ is the upper bound. Defining

$$\gamma_3 = \frac{A_+(1+\mu\tau_-) + \mu\tau_-\sqrt{A_+A_-}}{A_+(1+2\mu\tau_-) + (A_+ - A_-)\mu^2\tau_-^2}, \tag{4.36}$$

we require that $\gamma < \gamma_3$ in order that the segregated quasi-fixed points $\underline{s}_{QFP}$ may be stable. For $\gamma > \gamma_3$, the derivatives ds_j/dt, $\forall j \neq i$ become positive, so the segregated quasi-fixed points are certainly unstable in this region. It remains to be ensured that the nonzero, s_i direction is stable. At zeroth order, the various points $\underline{s}_{QFP}$ are part of the nonzero fixed hyperplane, which we know to be stable for $\gamma > \gamma_0$. Hence, the s_i direction is always stable at zeroth order. Thus, the segregated quasi-fixed points are certainly stable for $\gamma_0 < \gamma < \gamma_3$, where for consistency the lower bound γ_0 is of the same order in $\hat{\sigma}^2$, namely zeroth order, as the upper bound γ_3. This leaves open the small, first-order-sized window $\gamma_1 < \gamma < \gamma_0$ in which the segregated quasi-fixed points exist in the nonnegative hyperquadrant. In fact, the quasi-fixed points are stable in this small region too. Strictly speaking, it is inconsistent to write the condition on γ that guarantees the stability of the segregated quasi-fixed points in the form $\gamma_1 < \gamma < \gamma_3$, since the orders of the two bounds differ, but we shall do so anyway in order to close the small $\gamma_1 < \gamma < \gamma_0$ window.

We now turn to a real unsegregated fixed point. This fixed point is defined by $\underline{s}_{FP}$ having entirely nonzero components. Thus, since we are then forced to require

$$\left(s_+N_0 + \hat{\sigma}^2 N_1^i\right)|_{\underline{s}=\underline{s}_{FP}} = 0, \quad \forall i, \tag{4.37}$$

we must have $N_1^i = N_1^j$, $\forall i \neq j$. The simplest solution of this equation is $s_{FP}^i = s_{FP}^j$, $\forall i \neq j$, so that all the components of $\underline{s}_{FP}$ are equal (and nonzero). Because all the components are equal, we refer to this point as the *unsegregated fixed point*. We write $\underline{s}_{FP} = \frac{1}{m}(s_+, \ldots, s_+)^T$ and again expand s_+ as $s_+ = x_+ + \hat{\sigma}^2 y_+$. The zeroth-order solution must lie on the zeroth-order nonzero fixed hyperplane, so we again have

$$x_+ = \frac{\gamma(1+\mu\tau_-) - 1}{\mu\tau_+}. \tag{4.38}$$

We find that the first-order correction is

$$y_+ = \frac{\gamma}{\mu\tau_+}\left[\frac{1}{1+\mu\tau_-}\left(\mu\tau_- - \frac{1}{m}\gamma^{-2}\mu\tau_+x_+\right) - \frac{\mu\tau_-}{1+x_+}\left(1-\frac{1}{m}\right)\right]. \tag{4.39}$$

Again, we require that $s_+ = x_+ + \hat{\sigma}^2 y_+ > 0$ for this unsegregated fixed point to be accessible to the weight vector, so, expanding γ as $\gamma = \gamma_0 + \hat{\sigma}^2 \hat{\gamma}$ as for the segregated quasi-fixed points, we find that γ must satisfy

$$\gamma > \frac{1}{1+\mu\tau_-}\left[1 + \hat{\sigma}^2 \frac{\mu\tau_-}{1+\mu\tau_-}\left(\frac{\mu\tau_-}{1+\mu\tau_-} - \frac{1}{m}\right)\right]. \tag{4.40}$$

The right-hand side of this inequality is precisely γ_2 defined in equation 4.27, which determines the value of γ at which the zero fixed point turns from a saddle into a repeller. Thus, as the unsegregated fixed point passes through the origin into the positive hyperquadrant, the zero fixed point turns into a repeller. To determine the stability of the unsegregated fixed point, we expand and linearize about it as usual, and after lengthy algebra we find that

$$\frac{d\delta_i}{dt} = \frac{A_-\tau_-\mu}{1+\mu\tau_-}\frac{1}{(1+x_+)^2}\frac{1}{\gamma}\left[\frac{x_+(1+x_+)}{1+\mu\tau_-}\left(-\mu\tau_+ + \hat{\sigma}^2 J\right)\delta_+ - \hat{\sigma}^2 K\,\delta_i\right], \tag{4.41}$$

where J is a long and unwieldy expression that we do not reproduce here, and K is given by

$$K = x_+\left(\gamma - \frac{1}{1+\mu\tau_-}\right) - \gamma\frac{\mu\tau_-}{1+\mu\tau_-}. \tag{4.42}$$

The eigenvalues of the associated matrix are then just

$$\lambda_1 = \frac{mA_-\tau_-\mu}{1+\mu\tau_-}\frac{1}{(1+x_+)^2}\frac{1}{\gamma}\left[\frac{x_+(1+x_+)}{1+\mu\tau_-}\left(-\mu\tau_+ + \hat{\sigma}^2 J\right) - \frac{1}{m}\hat{\sigma}^2 K\right], \tag{4.43}$$

and

$$\lambda_i = -\frac{A_-\tau_-\mu}{1+\mu\tau_-}\frac{1}{(1+x_+)^2}\frac{1}{\gamma}\hat{\sigma}^2 K, \quad \forall\, i > 1. \tag{4.44}$$

Although the expression is messy, it is easy to show that λ_1, associated with the eigenvector $(1, \ldots, 1)^T$, changes sign at $\gamma = \gamma_2$, being positive for $\gamma < \gamma_2$ and negative for $\gamma > \gamma_2$. Hence, as the unsegregated fixed point moves into the positive hyperquadrant, the zero fixed point turns into a repeller because the direction corresponding to $(1, \ldots, 1)^T$ becomes unstable and the unsegregated fixed point becomes stable precisely in this same direction. The sign of all the other eigenvalues associated with the unsegregated fixed

point is determined solely by K. Stability in all the directions orthogonal to $(1, \ldots, 1)^T$ thus requires $K > 0$, or

$$x_+ \left(\gamma - \frac{1}{1 + \mu\tau_-} \right) - \gamma \frac{\mu\tau_-}{1 + \mu\tau_-} > 0. \tag{4.45}$$

This is identical to equation 4.33, determining the stabilities of the segregated quasi-fixed points, except that the inequality is opposite. Thus, we see immediately that we must have $\gamma > \gamma_3$ for the stability of the unsegregated fixed point, and for $\gamma_2 < \gamma < \gamma_3$, the unsegregated fixed point is an unstable saddle node. The unsegregated fixed point becomes stable precisely when the segregated quasi-fixed points become unstable, and the unsegregated fixed point becomes accessible to the weight vector precisely when the zero fixed point turns into a repeller. Thus, all three sets of fixed points are dynamically coupled in terms of their stabilities.

Because the local fixed-point structure must be globally consistent, we can deduce that there must exist other fixed points for $m \geq 3$. For example, in the interval $\gamma \in (\gamma_1, \gamma_3)$, there must exist saddles that partition the afferent weight vector space into m regions, each region being defined by the requirement that an initial weight vector in the region always flows to the same quasi-fixed point. We do not explore these additional, essentially uninteresting fixed points here because they merely render consistent the global fixed-point structure that is determined by the relative stabilities of the segregated quasi-fixed points and the unsegregated fixed point. It is the stabilities of the segregated and unsegregated states in which we are principally interested here, since these are the states relevant for a putative model of synaptic competition.

We see that the first-order dynamics are again uniquely determined by the value of γ, but in contrast to the zeroth-order dynamics, which had only two distinct regimes, the first-order system has four distinct regimes. We now summarize the dynamics in each of the regimes.

For $\gamma < \gamma_1$, there is only one fixed point, at the origin. Although a saddle node, it is stable in the $(1, \ldots, 1)^T$ direction. Hence, all flow initially moves parallel to this vector in the direction of the origin. If the weight vector moves sufficiently close to the origin, it will experience a repulsion in the directions orthogonal to $(1, \ldots, 1)^T$, but for $\gamma < \gamma_1$ this repulsion is never sufficiently strong to reverse the downward components of flow toward the origin. In all cases, the weight vector will eventually hit an $s_i = 0$, for some i, hyperplane and be trapped in it by the truncation procedure. Because there are still negative components of flow, the weight vector continues to move toward the origin, becoming trapped in further $s_j = 0$, $j \neq i$ hyperplanes. The weight vector thus is always driven toward the origin and ends up there, despite the origin's being a saddle node. The truncation at zero thus overrides this fixed point's instability.

For $\gamma_1 < \gamma < \gamma_2$, the zero fixed point is still a saddle node, and now there exists a set of m stable, segregated quasi-fixed points. In this regime of γ, the repulsion away from the origin is sufficiently strong to reverse the negative components of flow. Thus, for a weight vector sufficiently close to the origin, it hits an $s_i = 0$, for some i, hyperplane and is turned away from the origin, so that it starts moving in the opposite direction. It moves up $s_i = 0$ hyperplanes until it reaches a stable, quasi-fixed point and remains there. For weight vectors sufficiently distant from the origin, they flow toward the origin parallel to the direction $(1, \ldots, 1)^T$. These vectors never go sufficiently close to the origin to have their negative components of flow reversed. They thus hit $s_i = 0$ hyperplanes and move down toward the stable, quasi-fixed points, and stay there.

In the regime $\gamma_2 < \gamma < \gamma_3$, the zero fixed point has now turned into a repeller, so there are never any components of flow toward the origin in the neighborhood of the origin. An unsegregated fixed point has become accessible, but it is a saddle node. The segregated quasi-fixed points remain stable. Hence, this regime is essentially identical to the regime in which $\gamma_1 < \gamma < \gamma_2$, except with local differences near the origin and around the now-accessible unsegregated fixed point. All flow therefore ends up at the stable, segregated quasi-fixed points.

The combined regime $\gamma_1 < \gamma < \gamma_3$ therefore supports stable, competitive dynamics, allowing afferents to segregate on the target cell in an activity-dependent manner, as required, for example, in a model of ODC formation.

Finally, when $\gamma_3 < \gamma$, the segregated quasi-fixed points become unstable, and the unsegregated fixed point becomes an attractor. The origin remains a repeller. Hence, all flow ends up at the unsegregated fixed point.

The phase portraits for the two interesting, dynamically distinct γ regimes are shown in Figure 9 for a two-afferent system. We do not show the low γ regime, $\gamma < \gamma_1$, because the portrait is trivial, with all initial conditions flowing to the origin. Figure 9A shows the regime in which γ takes an intermediate value, $\gamma_1 < \gamma < \gamma_3$, for which a set of stable, segregated quasi-fixed points exists. The system always evolves to one of these segregated points. For this intermediate value of γ, the unsegregated fixed point is unstable. When γ is too high, $\gamma_3 < \gamma$, shown in Figure 9B, the segregated quasi-fixed points become unstable, and the unsegregated fixed point becomes stable, so afferent segregation on the target cell breaks down, and the afferent weight vector always flows to the unsegregated fixed point.

4.2 Full ∞-Spike Rule. The above fixed-point analysis for the simplified rule may be repeated for the full, nonresetting, ∞-spike rule in equation 2.22, for the convenient parameter choice $n_\pm = 1$. We do not present the results of this analysis here because the expressions that arise from the full model are unwieldy and thus lack the transparency of those for the simplified model. The simplified model has the virtue, compared to the full model, that almost all the resulting expressions can be stated on

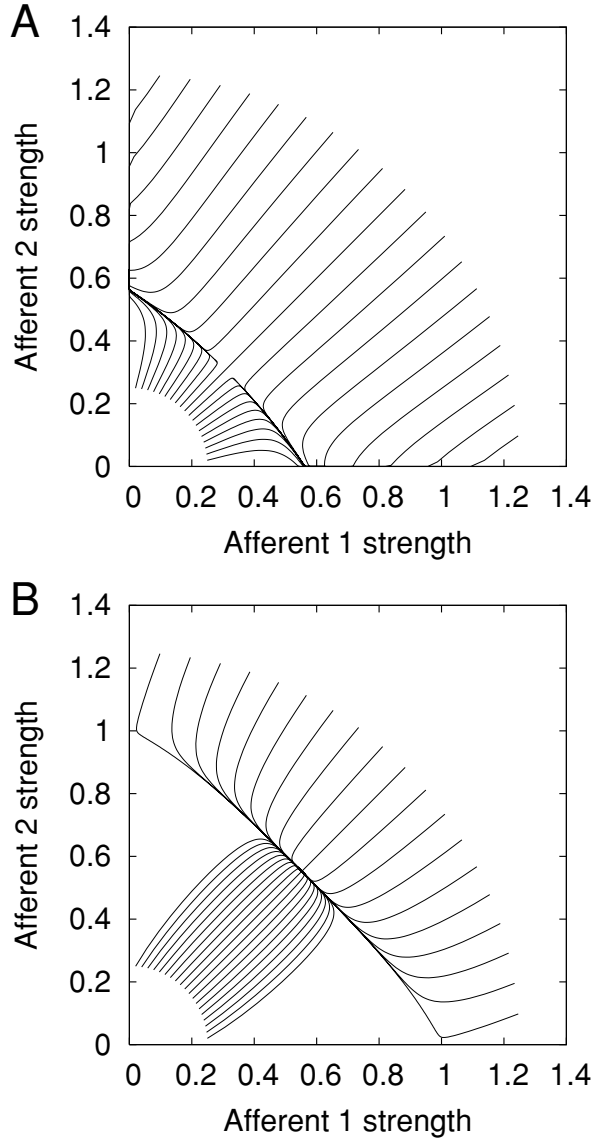

Figure 9: Phase portraits of the simplified, nonresetting ∞-spike rule, with $n_{\pm} = 1$, in a system with two afferents. (A) Evolution to the segregated quasi-fixed points with $\gamma = 0.65$. (B) Evolution to the unsegregated fixed point with $\gamma = 0.95$.

one line, while those for the full model occupy several lines. Nevertheless, the full model possesses dynamics that are qualitatively identical to those of the simplified model discussed above.

In the full model, we observe similar critical values of γ at which new quasi- or real fixed points become accessible to the afferent weight vector or at which the fixed points change their stability. In particular, corresponding to the values γ_1 and γ_2 for the simplified model, we have the values, say, $\overline{\gamma}_1$ and $\overline{\gamma}_2$, for the full model, where these differ from γ_1 and γ_2 by terms of order $\hat{\sigma}^2$. At $\gamma = \overline{\gamma}_1$, the segregated quasi-fixed points become available and are initially stable. At $\gamma = \overline{\gamma}_2$, the unsegregated fixed point moves into the nonnegative hyperquadrant, being initially an unstable saddle, and the zero fixed point at the origin simultaneously changes from a saddle node into a repelling node. We still have $\overline{\gamma}_1 \leq \overline{\gamma}_2$ as in the simplified model.

Corresponding to γ_3 in the simplified model, we also have a new value, say, $\overline{\gamma}_3$, in the full model. Although derived from first-order equations, γ_3 and $\overline{\gamma}_3$ are nonetheless purely zeroth-order in $\hat{\sigma}^2$. Hence, γ_3 and $\overline{\gamma}_3$ differ even at zeroth order. Despite this, the dynamics associated with the $\gamma = \overline{\gamma}_3$ transition in the full model are identical to those associated with the $\gamma = \gamma_3$ transition in the simplified model. At this point, the segregated quasi-fixed points become unstable, and the unsegregated fixed point simultaneously becomes stable.

Although our analyses of the full and simplified models have been performed only for the case $n_\pm = 1$, for reasons of analytical simplicity, we can explore numerically the impact of other values of $n_\pm$ on the full and simplified models. For $n_\pm = 3$, for example, we observe dynamics that are qualitatively similar to the $n_\pm = 1$ case, with the same three, essentially distinct parameter regimes in γ.

4.3 Beyond Small α_i. In the above fixed-point analysis, we expanded, for example, equation 4.6 in the variables α_i, representing small fluctuations about a common mean afferent activity, μ, and then performed an ensemble average over these fluctuations. Although this permits us to make some progress in terms of understanding the fixed-point dynamics of the model, it necessarily does not allow an examination of the large α_i regime, in which the fluctuations about the mean activity can be large. In principle, because we defined the size of these fluctuations with respect to the mean firing rate, so that $|\alpha_i| \leq 1$, we could continue the expansion to yet higher orders in $\hat{\sigma}^2$. Although possible, doing so would be tiresome and the resulting expressions an uncontrollable mess. It is therefore not clear that any additional analytical insight would be possible in the face of the growing complexity of the new terms.

We can, however, explore the fixed-point structure for large α_i by means of numerical simulation. When we do so, we find essentially the same three regimes in γ for both the full and simplified models considered above. In

particular, γ_1 and γ_2 still exist and define the points at which the segregated quasi-fixed points and the unsegregated fixed point, respectively, move into the nonnegative hyperquadrant and thus become accessible to the weight vector. However, the γ_3 critical value in both models is somewhat modified. The large α_i fluctuations "split" this value of γ into two different values; call them, say, γ_3' and γ_3'', where $\gamma_3' < \gamma_3''$. At $\gamma = \gamma_3'$, the unsegregated fixed point becomes stable, while the segregated quasi-fixed points remain stable. Only at $\gamma = \gamma_3''$ do the segregated quasi-fixed points become unstable. Thus, there is a narrow interval for γ, $\gamma \in (\gamma_3', \gamma_3'')$, of size of order $\hat{\sigma}^2$ or higher, in which the system may evolve to either a segregated quasi-fixed point or the unsegregated fixed point. To which point the system evolves is determined by the initial conditions. Because the fixed-point structure must be globally consistent, we deduce that there must exist in this narrow regime a new set of saddle fixed points that partition the space into a region containing the unsegregated fixed point, to which any initial point in this region will flow, and m regions containing the segregated quasi-fixed points. The existence of this narrow transition region in which both the segregated quasi-fixed points and the unsegregated fixed point are simultaneously stable is all that appears to distinguish the dynamics of both the full and simplified models in the small α_i, analytically explored regime from the large α_i, numerically explored regime. It is likely, in fact, that this transition region is present even for small α_i, but is so narrow as to be extremely difficult to observe numerically.

5 Computation in the Rate-Based Limit

To derive the n-spike, rate-based rules, we integrated over the interspike intervals and averaged over all 2^n possible spike trains to compute an unconditional expectation value for the change in synaptic efficacy due to a typical n-spike train. The resulting rate-based rule is somewhat abstract, in the sense that the neuron does not really compute at the level of this rate-based rule but continues to compute at the level of individual spikes. We may ask, however, whether there are any conditions under which, although the neuron is computing at the level of spikes, it nevertheless behaves as if it were following the rate-based, n-spike rule. If no such conditions exist, then our analysis is somewhat academic, since the derived rules would be merely mathematical abstractions to which the neuron's behavior cannot ever approximate.

We can identify two limiting cases of interest. First, if the spike train is extremely long and not highly unusual, then we can think of the train as naturally decomposing into a set of shorter subtrains. The neuron can then effectively average its behavior over these subtrains. If there are enough subtrains, then sufficient averaging over them will occur in order for the neuron's mean dynamics to approximate the rate-based rules. However, the

variance in this behavior could be large, and these fluctuations could thus nevertheless prevent the emergence of stable, segregated states at the spike-based level. Second, we know that even for small n, $n \geq 3$, the rate-based rules exhibit stable, segregated states, and the above considerations for large n do not apply. Of course, if we present a neuron with a single instance of an n-spike train, n small, then we would not expect its synaptic strengths to evolve much during this short train. The neuron must therefore be presented with a long sequence of such n-spike trains, so that it can average over this sequence, and the trains must be sufficiently well separated that all the synapses return to the *OFF* state between trains. Again, the averaging will ensure that the mean behavior is exhibited, but the fluctuations could destroy the stability. We therefore see that the key to stability is to ensure that the fluctuations are small.

Small fluctuations can be guaranteed provided that the neuron's dynamics are not dominated by the most recent spike train (or subtrain). If the most recent spike train essentially erases the neuron's state developed from exposure to earlier trains, then the neuron's behavior will be dominated by train-to-train fluctuations. We therefore require that the change in synaptic strength induced by each spike train is small compared to the (nonzero) synaptic strengths. This can be achieved by setting $A_{\pm}$ sufficiently small, since these two parameters set the overall magnitude of plasticity.

Changing the magnitude of plasticity (the overall scale of $A_{\pm}$) is, of course, equivalent to changing the learning rate in a model's dynamics. The learning rate is essentially the step size in a numerical integration procedure. It is well known that the stability of the numerical solutions of a set of differential equations depends critically on the step size and that there usually exists a threshold above which the integration scheme fails to converge to the exact solutions, with chaotic instability or divergent behavior ensuing. We should therefore not be too surprised that when the magnitude of plasticity is sufficiently small, the spike-based behavior will be expected to converge to the rate-based behavior. Given the complexity of the switch model, however, determining the location of the threshold above which the neuron does not compute in the rate-based limit, and exhibits instead a strong dependence primarily on the last spike train, is a difficult matter. We therefore resort to a simple numerical search for the approximate location of this threshold.

To obtain a condition on the magnitude of plasticity below which the rate-based behavior becomes dominant in a spike-based simulation, we determine when the spike-based system exhibits qualitatively the same fixed-point structure known to exist in the rate-based system. We consider a system of two afferents for simplicity. A rate-based simulation of two afferents will stably segregate, with one afferent gaining complete control of the target cell, provided that $\gamma_1 < \gamma < \gamma_3$. Thus, we select a value of γ in this range and run a spike-based simulation for various values of the overall scale of $A_{\pm}$. In these simulations, in order to examine the n-spike rule, we

present a series of n-spike trains to each of the afferent's synapses, and after every train force the synapses to return to the *OFF* state, which is equivalent to spacing the trains sufficiently far apart that they do not interact. Within each train, the afferents have their Poisson firing rates randomly fixed either "high" (75 Hz) or "low" (25 Hz). We perform an initial presentation of 2.5×10^7 spikes, partitioned into n-spike trains, in order to allow sufficient time for the afferents to segregate on the target cell. At a typical, average rate of 50 Hz, this corresponds to approximately six days' worth of simulated synaptic activity, which is not too dissimilar to the typical timescale for developmental processes in the nervous system. After this initial period to allow time for segregation, we present another series of 2.5×10^7 spikes, again partitioned into n-spike trains, during which we probe the extent and stability of any segregation. After each train presentation, we calculate the segregation index, S_I, which we define as

$$S_I = \frac{s_1 - s_2}{s_1 + s_2}. \tag{5.1}$$

If the afferents are well segregated, then S_I will take values close to $+1$ or -1, depending on which afferent controls the target cell. If segregation is stable, then this index will not change much, except for small fluctuations. If the afferents are segregated but not stably so, with control switching between the two afferents, then S_I will flip between $+1$ and -1. Averaged over sufficient trains, its value will be roughly zero. If the afferents are not segregated, but oscillate about a mean synaptic strength, then S_I will always be roughly zero, and of course its average will be roughly zero. Thus, for this probing phase, we determine $\langle S_I \rangle_P$, where $\langle \rangle_P$ denotes the average value of S_I during this second period. We then take the absolute value of this average, and average this value over 50 distinct runs for each value of the overall scale of $A_\pm$. Thus, our final measure of segregation and stability is $\langle |\langle S_I \rangle_P| \rangle_R$, where $\langle \rangle_R$ denotes an average over runs.

In Figure 10, we plot $\langle |\langle S_I \rangle_P| \rangle_R$ as a function of the overall scale of $A_\pm$ for 3-, 9- and 15-spike train simulations for the nonresetting model with $n_\pm = 3$ and $\gamma = 0.6$. We obtain qualitatively similar results for the resetting model and for different values of γ. Also shown in Figure 10 is the fit of our raw data to logistic-like functions, $a - b \tanh(cx - d)$, where a, b, c, and d are fitted parameters. That the fits to logistic-like functions match the raw data well indicates that the transition from stable segregation to unstable segregation is relatively sharp. For an overall scale of plasticity of approximately 10^{-3} or lower, depending on the number of spikes in the train, we observe robust and stable afferent segregation, while for a scale greater than this value, segregation is achieved, but is not stable, so that the afferents change their control of the target cell over time. For values of the overall scale very much greater than 10^{-3}, segregation is not achieved at all. We observe that as the train length increases, the mean segregation

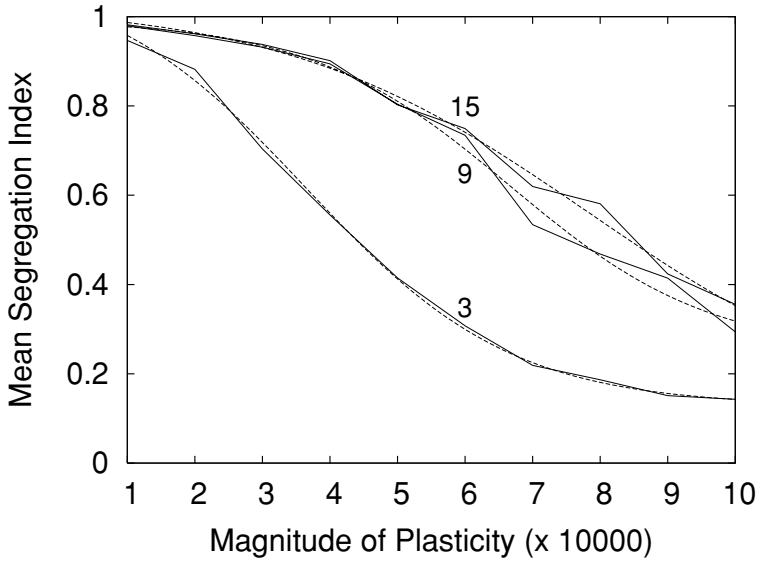

Figure 10: The dependence of the mean segregation index $\langle|\langle S_I\rangle_P|\rangle_R$ on the magnitude of plasticity, or the overall scale of $A_\pm$. Solid lines represent the numerically obtained values of $\langle|\langle S_I\rangle_P|\rangle_R$ for the number of spikes in each train indicated by the attached number. The dashed lines show the fit of the raw data to logistic-like functions, as described in the text.

index increases for a fixed value of the overall plasticity magnitude. However, by about 10 to 20 spikes, asymptotic behavior is reached, with no further increase in the index observed. This is in accord with our expectations, since the n-spike rate-based learning rules converge very rapidly as a function of n, with convergence achieved by $n \approx 10$ spikes (Appleby & Elliott, 2006).

The data in Figure 10 are obtained by randomly fixing the afferents' Poisson firing rates within each spike train. For longer trains, therefore, each afferent fires for longer with the same firing rate. It could therefore be argued that the dependence of the magnitude of plasticity on the number of spikes in a train merely reflects this longer exposure to the same firing pattern. We can rule this out in two ways. First, instead of fixing each train's afferents' rates, we can instead fix each afferent's rate for a given period of time (a firing "epoch"), regardless of the number of spikes in each train. For a firing epoch length of 1000 ms, for example, we obtain data that are essentially identical to those in Figure 10 (data not shown). Second, we can, for example, consider three-spike trains with fixed firing per train but simply duplicate the firing patterns between consecutive trains, so that

two-train sequences of three-spike trains have the same rates. Doing this, we produce data identical to the three-spike data shown in Figure 10 and not the six-spike data with an enhancement in the segregation index (data not shown). Thus, longer exposure to the same activity patterns is not responsible for the trend exhibited in Figure 10. Rather, the observed trend reflects the fact that as the number of spikes increases in a train, competition becomes stronger and stronger, as revealed by equation 3.22.

A modification level of 10^{-3} roughly translates to a 0.1% change in synaptic strength per potentiation or depression event. For an afferent supporting 10 synapses, this translates to a 1% change in afferent strength per spike pairing event. This approximate magnitude is commensurate with that seen in experimental work, where a change of around 1.5% per spike pairing is typically measured (Bi & Poo, 1998).

Although this numerical treatment is not exact, it is sufficient for our purposes to demonstrate that computation in the rate-based limit is realistically available to a real, spike-based system and to have an approximate idea of where that limit lies. Assuming that $A_{\pm}$ satisfy the numerically obtained condition on the magnitude of plasticity below which the rate-based behavior becomes dominant, we may dispense with spikes completely and work in the rate-based limit. The analysis of section 4 was carried out in the rate-based limit and did not consider individual spikes. The stable, competitive dynamics that we observed will therefore be the dominant mode of computation provided that $A_{\pm} \sim 10^{-3}$ or lower.

6 Large-Scale Numerical Simulations

In section 4 we established that the ∞-spike nonresetting rule possesses a fixed-point structure consistent with the presence of stable, competitive dynamics in the model. The presence of such dynamics has also been argued for in all multispike rules for an appropriate choice of parameters. In section 5, we established that the rate-based multispike dynamics are actually available to neurons in the sense that, although computing at the level of single spikes, the neuron's dynamics will respect the fixed-point structure of the rate-based rules and exhibit stable, competitive dynamics. We are therefore now in a position to consider a large-scale simulation of, for example, ODC development, in order to demonstrate that the results shown for a single target cell scale up without difficulty to multiple target cells. Moreover, while for convenience we considered above only uncorrelated afferent activity patterns, a simulation of ODC development will permit us to explore negatively or positively correlated afferent activity patterns too. For ODC development in the cat, which develops ODCs in the presence of presumably positively correlated afferent activity after eye opening, the case of positive correlations is more pertinent. This is also usually the much harder task: models that can segregate negatively correlated afferents

frequently cannot segregate positively correlated afferents. For our switch model to be a candidate model of developmental synaptic plasticity, we must therefore show that it operates successfully in the face of positively correlated afferent activity patterns.

We run ODC simulations according to existing, documented protocols (Elliott & Shadbolt, 1998). We use the rate-based, nonresetting ∞-spike rule with $n_{\pm} = 1$, in equation 2.22, as the synaptic modification rule. In brief, we simulate two patches of retinotopically equivalent lateral geniculate nucleus (LGN), each containing a square array of cells of size 13×13, with periodic boundary conditions imposed for convenience. The cortex is a 25×25 square array of cells, again with periodic boundary conditions. Each LGN cell arborizes over a retinotopically appropriate patch of cortex, of size 7×7. LGN activity patterns are constructed according to the method of Goodhill (1993), taking the form of gaussian correlated noise. The parameter $p \in [0, 1]$ determines the activity correlations between the two LGN patches, with $p = 0$ corresponding to perfectly anticorrelated patterns and $p = 1$ corresponding to perfectly correlated patterns. Cortical activity is just the standard linear sum of afferent input, but smeared by convolving cortical activity with a short-range gaussian function. Such smearing could be achieved, for example, by lateral excitation. In this way, nearby cortical cells fire similarly, which is necessary in order to develop structured ODCs rather than a pattern of salt-and-pepper segregation. It is well known that presynaptic constraints must typically be introduced in order for the pattern of ODCs to exhibit some degree of regularity, with fairly constant stripe widths across the whole simulated patch of cortex. To achieve this, we introduce a lower bound on the total synaptic strength that an afferent may support. If it falls below 75% of the average total strength supported by all afferents, we simply freeze depression. We use this method because it is convenient and removes the need for us to calculate explicitly the expected total afferent strength in such a simulation and then set a lower bound accordingly.

In Figure 11 we show three ODC maps corresponding to three different values of p: $p = 0.3$ representing negatively correlated activity patterns; $p = 0.5$ representing uncorrelated activity patterns; and $p = 0.7$ representing positively correlated activity patterns. We see clear patterns of ODCs in all cases, with well-segregated afferents, although with increasingly binocular boundaries between ODCs as p increases, as expected. We also observe a clear decrease in the widths of ODCs as the inter-ocular correlations increase. This phenomenon was first observed by Goodhill (1993) in simulation, and subsequent experimental results supported the possibility that the widths of ODCs are not fixed but may be partially determined by visual experience (Löwel, 1994; Tieman & Tumosa, 1997). Our results here show that our switch model of STDP scales up to a large-scale simulation without difficulty and can comfortably segregate afferents whose activities are positively correlated.

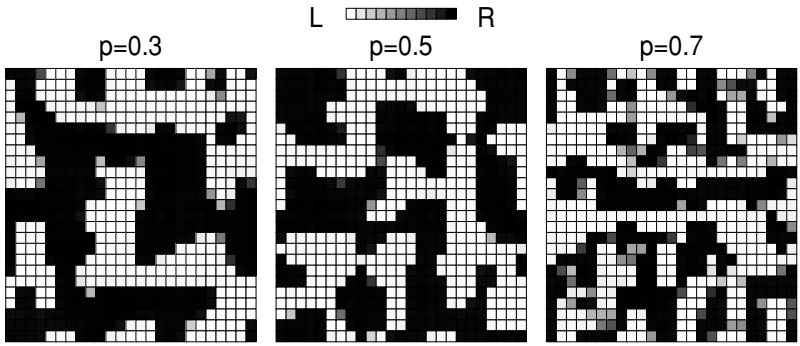

Figure 11: A simulation of ocular dominance column formation in the non-resetting ∞-spike model with $n_{\pm} = 1$ for the three values of the interocular correlation probability shown above each map. In these maps, each square represents a single cortical neuron. The shade of gray assigned indicates the relative degree of control by the two eyes. A white square indicates complete control by the left eye, and a black square indicates complete control by the right eye. Shade of gray, as shown in the key, interpolate between these two extremes.

7 Discussion

In this letter, we have studied the synaptic dynamics induced by the two- and multispike rules that are derived within the context of a stochastic model of STDP. In this model, the STDP rule emerges only at the temporal and synaptic ensemble level, while individual synapses obey a much simpler, computationally less burdensome plasticity rule in which synaptic strengths change in an all-or-none fashion independent of spike timing. The structure of the postulated unified three-state synaptic switch, which was initially designed only to account for the basic phenomenology of the STDP learning rule seen in the context of interactions between two spikes, tightly constrains the form of the multispike interactions that exist in our model. One freedom that the model possesses is to allow stochastic process resetting in the *DEP* and *POT* states. Although resetting changes the precise form of the multispike interaction functions, nevertheless, dynamically speaking, the qualitative dynamics observed, in terms of the fixed-point structure of the model, are essentially independent of whether we allow resetting.

Without further modification, we find that the two-spike learning rule is irredeemably pathological in its learning dynamics, sending all afferents to either zero strength or unbounded growth. The two-spike learning rule cannot therefore support the stable, competitive dynamics that are essential within the context of developmental synaptic plasticity. This observation is particularly interesting given that the majority of theoretical studies of

STDP are based on two-spike rules. Often the introduction of additional constraints and nonlinearities becomes necessary, such as imposing hard upper bounds on synaptic strengths (Song et al., 2000), scaling of potentiation with synaptic strength (van Rossum et al., 2000), or the temporal restriction of spike interactions (Izhikevich & Desai, 2003). These various additional constraints are required to stabilize the underlying STDP learning dynamics of the unmodified STDP rule. The assertion that the two-spike STDP rule can therefore account for, for example, developmental synaptic plasticity must be construed with caution. It appears, in fact, that the additional constraints required to stabilize the two-spike STDP rule's dynamics are actually playing a significant role in giving rise to the very dynamics of interest.

The BCM model (Bienenstock et al., 1982) is a well-studied model of developmental synaptic plasticity, exhibiting the stable, competitive dynamics required of such a model. Systematic attempts have therefore been made to connect the two-spike STDP rule to the BCM model, in order to endow the STDP rule with the dynamics exhibited by the BCM rule (Izhikevich & Desai, 2003). In our own formulation of STDP, we find that a large-parameter regime exists in which the rate-based, two-spike learning rule is qualitatively BCM-like. It is, however, exactly this BCM-like form that gives rise to the two-spike rule's instabilities, trapping afferents in a depressing well for small synaptic strengths and permitting unbounded growth for large synaptic strengths. This behavior occurs because the threshold determining the cross-over point between depression and potentiation in our unmodified two-spike rule is fixed, and, indeed, the BCM model with such a fixed rather than a sliding threshold would exhibit identical dynamics. By introducing such a sliding threshold, the two-spike rule can therefore be stabilized, as expected.

Focusing attention only on two-spike interactions and thus attempting to patch up the dynamics exhibited by the two-spike STDP learning rule strikes us as unnecessarily restrictive. After all, a real neuron does not experience a disconnected series of spike pairs, but rather an entire train of freely interlacing pre- and postsynaptic spike events. The dynamics of any plasticity rule should thus be understood in the context of these natural spike trains, and not in the context of an experimental stimulation protocol that happens to be particularly convenient to the experimentalist. It is therefore critical that multispike interactions be understood, especially given the difficulties arising from the two-spike interaction function. However, it is precisely the generalization to longer spike trains that creates difficulties for many existing models of STDP. Such devices as making the scale of plasticity sensitive to spike history (Froemke & Dan, 2002) or restricting the temporal interactions between spikes (Izhikevich & Desai, 2003) are introduced. To be sure, such methods do replicate spike triplet data (Froemke & Dan, 2002) and do allow a BCM rule to be derived from STDP rules (Izhikevich & Desai, 2003). However, such approaches seem to us essentially to be attempts to

force multispike interactions into the strait jacket imposed by two-spike interactions.

Without prior design (since we were thinking only about two-spike STDP), our model essentially forces on us all the multispike interaction functions, and all such functions exhibit qualitatively similar dynamics—dynamics that differ significantly from those exhibited by the two-spike interaction function. For a very broad range of parameters, all these rules exhibit stable, competitive dynamics without any further modification. Even when the lower-order multispike rules fail, we are guaranteed that all the higher-order multispike rules operate as required. At first blush, this is astonishing. Merely extending our considerations from two spikes to three spikes immediately solves all the problems inherent in the two-spike rule.

Why should the step from two to three spikes change the dynamics so dramatically? We argued that this is because in the unified switch rule, depression and potentiation are coupled. By this rather vague notion, we mean that activation of the potentiation lobe of the switch precludes a simultaneous activation of the depression lobe of the switch, and vice versa. In a sense, a putative potentiation event blocks a possible depression event, and vice versa. By breaking the two lobes into two separate switches, so that they may be simultaneously active, we showed that competition immediately breaks downs, resulting in the runaway learning characteristic of the two-spike rule. Critical to probing this coupling between potentiation and depression is the presence of three or more spikes; even in the unified switch, two spikes can never probe this couplings, and so two-spike interactions always induce uncoupled potentiation and depression dynamics. It therefore appears that this coupling, or interaction, between potentiation and depression processes at the level of a single neuron is vital to the presence of stable, competitive dynamics. Indeed, a sliding threshold in the BCM model provides precisely a means of coupling potentiation and depression events, in that the threshold is sensitive to the firing history of the cell, and this firing history is dependent on the strengthening and weakening of the synapses that the cell supports. Without this coupling, so that the threshold is fixed, stability and competition in the BCM model break down.

It appears in general, therefore, that in order for a model of synaptic competition to operate successfully, the machinery of potentiation and the machinery of depression cannot be independent. Although the dependence of the two processes is very indirect in the BCM model, via the sliding threshold, it is nevertheless critical to the model's successful operation. We therefore speculate that any model in which potentiation and depression are completely independent processes cannot be competitive. It would be interesting to seek to prove this within the context of the most general class of model possible, although this is likely to be very difficult.

In the above analyses, we have truncated s_i at zero when it is driven negative. We saw, for example, that although the zero fixed point is a saddle

for $\gamma < \gamma_2$ in the simplified ∞-spike resetting model, nevertheless, the truncation procedure dynamically converts it into a stable node when $\gamma < \gamma_1$. Moreover, the truncation gave rise, partly, to the existence of quasi-fixed points. It might therefore appear that this truncation procedure is playing an important role in our model's dynamics, being partly responsible for the presence of stable, competitive dynamics in it. This is not, however, the case. We can see this in two different ways.

First, from a mathematical point of view, we set $ds_i/dt = \Delta S_n(\lambda_{\pi_i}, \lambda_p)$ and used this as our synaptic strength update rule. This is not the only choice available to us. In this formulation, s_i represents the average synaptic strength over all of afferent i's synapses on the target cell. Suppose, however, that we instead think of s_i as the (possibly scaled) number of synapses supported by afferent i, each of which experiences the single-synapse learning rule. Then the overall change in afferent i's synaptic strength will be the individual change multiplied by the number of synapses, s_i. Hence, on this view, we would instead write $ds_i/dt = s_i \Delta S_n(\lambda_{\pi_i}, \lambda_p)$. The presence of the extra factor of s_i in this rule immediately means that synapses cannot evolve below zero, and thus the quasi-fixed points are converted into real fixed points, with their location and stability completely unchanged. Furthermore, it is easy to see that the zero fixed point in this formulation is genuinely stable for $\gamma < \gamma_1$. We further note that truncation at zero is a hard nonlinearity to which the afferents are completely insensitive when $s_i > 0$. Thus, their evolution to the vicinity of the segregated quasi-fixed points in the interval $\gamma \in (\gamma_1, \gamma_3)$ is entirely independent of truncation: it is a dynamical consequence of the overall structure of the model. The need for truncation, then, is merely a technical issue that does not fundamentally modify the model's prior dynamics.

Second, from a biological point of view, when a synapse reaches zero or close-to-zero synaptic strength, we expect it either to shut down and evolve no further, or to be retracted entirely. Biologically, an excitatory synapse of zero strength cannot turn into an inhibitory synapse. Our derivation of the learning rules, however, takes no account of this, for the purposes of tractability. However, the simplest remedy to this situation is merely to prevent a synapse from depressing when such depression would takes its synaptic strength negative. (Indeed, this it the presynaptic constraint that we used to model ODC development, in order not to obtain irregular patterns of afferent segregation, except that the lower limit on synaptic strength was set at some nonzero value.) This prevention of depression is tantamount to switching off the depressing lobe of the switch, perhaps by setting $A_- = 0$, when the synapse is too weak to permit further depression. In such a state, it could be potentiated but not depressed. This, though, is essentially equivalent to truncation of strengths at zero. The key competitive dynamics in the model are those that precisely allow a synapse to move toward a quasi-fixed point near which truncation, or the closing down of depression, then becomes necessary.

We observe three essentially distinct parameter regimes in our model's dynamics, corresponding to different ranges for the parameter γ. For low γ, all afferent strengths fall to zero. For high γ, afferents evolve to an unsegregated fixed point in which all afferents equally control the target cell. Only for intermediate values of γ do we see the evolution of the system to stable, segregated final states. Considered within the context of a simulation of ODC formation, the low- and high-γ regimes would correspond to a breakdown of ODC formation, while the intermediate regime would correspond to the normal development of ODCs. Do these three regimes correspond to real experimental situations, or are they just mathematical fictions? Interestingly, the infusion of the neurotrophic factors brain-derived neurotrophic factor (BDNF) or NT-4/5 (Cabelli, Hohn, & Shatz, 1995) or the blockade of their common, endogenous receptor trk-B (Cabelli, Shelton, Segal, & Shatz, 1997) both result in abolishing ODC development. In the former case, autoradiographic labeling reveals a higher-than-normal labeling, while in the latter case, the labeling is lower than normal. One interpretation of these results is that BDNF or NT-4/5 infusion causes a growth of afferent axonal arbors, while removing available BDNF or NT-4/5 from afferents causes their axonal arbors to atrophy. A similar influence of neurotrophic factors on axonal branching is also observed in the frog retinotectal system (Cohen-Cory & Fraser, 1995). Low and high γ, which cause the breakdown of ODC development, could therefore correspond to these experimental regimes. Moreover, low γ corresponds to very weak synapses, while high γ corresponds to strong synapses. Under an anatomical interpretation of synaptic strength, these would correspond to small and large axonal arbors, respectively. Intermediate values of γ then correspond to normal patterns of development. Given the capacity of neurotrophic factors rapidly to modulate synaptic transmission in the visual system (Akaneya, Tsumoto, & Hatanaka, 1996; Carmignoto, Pizzorusso, Tia, & Vicini, 1997; Sala et al., 1998), we can thus imagine a scenario in which neurotrophic factors dynamically determine the parameter γ in our model of STDP.

In conclusion, we have extended our earlier analysis of a stochastic model of STDP (Appleby & Elliott, 2005) to include an examination of multispike interactions. We have found that a consideration of multispike interactions is sufficient to endow our model with a fixed-point structure consistent with the presence of stable, competitive dynamics. In contrast, at least in our formulation, two-spike interactions by themselves cannot give rise to competitive dynamics. Multispike interactions therefore appear critical to understanding the presence of stable, competitive dynamics under STDP.

Acknowledgments

P.A.A. thanks the University of Southampton for the support of a studentship.

References

Akaneya, Y., Tsumoto, T., & Hatanaka, H. (1996). Brain-derived neurotrophic factor blocks long-term depression in the rat visual cortex. *J. Neurophysiol.*, *76*, 4198–4201.

Appleby, P. A., & Elliott, T. (2005). Synaptic and temporal ensemble interpretation of spike timing dependent plasticity. *Neural Comput.*, *17*, 2316–2336.

Appleby, P. A., & Elliott, T. (2006). Multi-spike interactions in a stochastic model of spike timing dependent plasticity. *Neural Comput.* (In press.)

Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.*, *18*, 10464–10472.

Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, *2*, 32–48.

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

Cabelli, R. J., Hohn, A., & Shatz, C. J. (1995). Inhibition of ocular dominance column formation by infusion of NT-4/5 or BDNF. *Science*, *267*, 1662–1666.

Cabelli, R. J., Shelton, D. L., Segal, R. A., & Shatz, C. J. (1997). Blockade of endogenous ligands of trkB inhibits formation of ocular dominance columns. *Neuron*, *19*, 63–76.

Carmignoto, G., Pizzorusso, T., Tia, S., & Vicini, S. (1997). Brain-derived neurotrophic factor and nerve growth factor potentiate excitatory synaptic transmission in the rat visual cortex. *J. Physiol.*, *498*, 153–164.

Castellani, G. C., Quinlan, E. M., Cooper, L. N., & Shouval, H. Z. (2001). A biophysical model of bidirectional synaptic plasticity: Dependence on AMPA and NMDA receptors. *Proc. Natl. Acad. Sci. USA*, *98*, 12772–12777.

Cohen-Cory, S., & Fraser, S. E. (1995). Effects of brain-derived neurotrophic factor on optic axon branching and remodelling in vivo. *Nature*, *378*, 192–196.

Dudek, S. M., & Bear, M. F. (1992). Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proc. Natl. Acad. Sci. USA*, *89*, 4363–4367.

Elliott, T., & Shadbolt, N. R. (1998) Competition for neurotrophic factors: Ocular dominance columns. *J. Neurosci.*, *18*, 5850–5858.

Froemke, R. C., & Dan, Y. (2002). Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, *416*, 433–438.

Goodhill, G. J. (1993). Topography and ocular dominance: A model exploring positive correlations. *Biol. Cybern.*, *69*, 109–118.

Gustafsson, B., Wigström, H., Abraham, W. C., & Huang, Y. -Y. (1987). Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus to single volley synaptic potentials. *J. Neurosci.*, *7*, 774–780.

Izhikevich, E. M., & Desai, N. S. (2003). Relating STDP to BCM. *Neural Comput.*, *15*, 1511–1523.

Karmarkar, U. R., & Buonomano, D. V. (2002). A model of spike-timing dependent plasticity: One or two coincidence detectors? *J. Neurophysiol.*, *88*, 507–513.

Löwel, S. (1994). Ocular dominance column development: Strabismus changes the spacing of adjacent columns in cat visual cortex. *J. Neurosci.*, *14*, 7451–7468.

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.*, *102*, 9679–9684.

Peterson, 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. USA*, *95*, 4732–4737.

Purves, D., & Lichtman, J. W. (1985). *Principles of neural development*. Sunderland, MA: Sinauer.

Roberts, P. D., & Bell, C. C. (2002). Spike timing dependent synaptic plasticity in biological systems. *Biol. Cybern.*, *87*, 392–403.

Sala, R., Viegi, A., Rossi, F. M., Pizzorusso, T., Bonanno, G., Raiteri, M., & Maffei, L. (1998). Nerve growth factor and brain-derived neurotrophic factor increase neurotransmitter release in the rat visual cortex. *Europ. J. Neurosci.*, *10*, 2185–2191.

Schuett, S., Bonhoeffer, T., & Hübener, M. (2001). Pairing-induced changes of orientation maps in cat visual cortex. *Neuron*, *32*, 325–337.

Shouval, H. Z., Bear, M. F., & Cooper, L. N. (2002). A unified model of NMDA receptor-dependent bidirectional synaptic plasticity. *Proc. Natl. Acad. Sci. USA*, *99*, 10831–10836.

Sjöström, P. J., Turrigiano, G. G., & Nelson, S. B. (2001). Rate, timing and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, *32*, 1149–1164.

Song, S., Miller, K., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat. Neurosci.*, *3*, 919–926.

Tieman, S. B., & Tumosa, N. (1997). Alternating monocular exposure increases the spacing of ocularity domains in area 17 of cats. *Visual Neurosci.*, *14*, 929–938.

van Ooyen, A. (2003). *Modeling neural development*. Cambridge, MA: MIT Press.

van Rossum, M. C. W., Bi, G. Q., & Turrigiano, G. G. (2000). Stable Hebbian learning from spike timing-dependent plasticity. *J. Neurosci.*, *20*, 8812–8821.

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., & Poo, M. M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, *395*, 37–44.

Received October 21, 2005; accepted March 27, 2006.