Competition for a limited supply of synaptic building blocks predicts multiplicative synaptic normalization and heterosynaptic plasticity

Jochen Triesch[1,*], Anne-Sophie Hafner[2]

1 Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany

2 Max-Planck Institute for Brain Research, Frankfurt am Main, Germany

∗ **E-mail: triesch@fias.uni-frankfurt.de**

Abstract

We present a mathematical model of synaptic normalization and heterosynaptic plasticity based on competition for limited synaptic resources. In the model, afferent synapses on a part of the dendritic tree of a neuron compete for a limited supply of synaptic building blocks such as AMPA receptors or other postsynaptic components, which are distributed across the dendritic tree. These building blocks form a pool of parts that are ready for incorporation into synapses. Using minimal assumptions, the model produces fast multiplicative normalization behavior and leads to a homeostatic form of heterosynaptic plasticity. It therefore supports the use of such rules in neural network models. Furthermore, the model predicts that the amount of heterosynaptic plasticity is small when many building blocks are available in the pool. The model also suggests that local production and/or assembly of postsynaptic building blocks across the dendritic tree may be necessary to maintain a neuron's proper function, because it facilitates their homogeneous distribution across the dendritic tree. Because of its simplicity and analytical tractability, the model provides a convenient starting point for the development of more detailed models of the molecular mechanisms underlying different forms of synaptic plasticity.

Author Summary

Changes in the efficacies of synapses are thought to be the neurobiological basis of learning and memory. When a synapse is strengthened, new neurotransmitter receptors are added to the postsynaptic membrane. Recent experiments have shown that the behavior of these receptors is highly dynamic, with receptors moving back and forth between synapses on time scales of seconds and minutes, reflecting a competition of synapses for available receptors. Here we propose a mathematical model of this competition of synapses for neurotransmitter receptors or other synaptic building blocks. Using minimal assumptions the model produces a multiplicative normalization behavior of synapses and it predicts a well-known form of so-called heterosynaptic plasticity, where changes in stimulated synaptic pathways induce changes of opposite sign in neighboring, unstimulated pathways. Thus, the model offers a parsimonious mechanistic explanation of these forms of synaptic plasticity.

Introduction

Simple mathematical models of Hebbian learning exhibit an unconstrained growth of synaptic efficacies. To avoid runaway dynamics, some mechanism for limiting weight growth needs to

be present. There is a long tradition of addressing this problem in neural network models using synaptic normalization rules [14, 20, 21, 23, 36]. Obviously, in order to keep up with the pace of synaptic changes due to Hebbian plasticity, normalization mechanisms must act sufficiently fast. Slow homeostatic synaptic scaling mechanisms [34] may therefore be ill-suited for ensuring stability [4, 36, 37]. A particularly interesting fast normalization rule scales synapses multiplicatively such that the sum of synaptic weights remains constant. Attractive features of such a rule, next to its conceptual simplicity, are that the relative strength of synapses are maintained and that in combination with Hebbian mechanisms it naturally gives rise to lognormal-like weight distributions as observed experimentally [17, 22, 28, 38]. While such normalization mechanisms are not considered biologically implausible, their link to neurobiological experiments has been tenuous, however.

In a recent review, Chiastiakova and colleagues argue that so-called heterosynaptic plasticity [2, 19] may be a prime candidate for such a fast synaptic normalization scheme [4]. The term "heterosynaptic" plasticity is used in contrast to the much more widely studied "homosynaptic" plasticity, where changes occur in a stimulated synaptic pathway. In contrast, heterosynaptic plasticity refers to changes in synaptic efficacies that occur in an unstimulated pathway after the stimulation of a neighboring pathway. The most common form of heterosynaptic plasticity has a homeostatic nature: If synapses in stimulated pathways potentiate, then this is accompanied by a depression of unstimulated pathways. Conversely, if synapses in stimulated pathways depress, this is accompanied by a potentiation of unstimulated pathways. A classic example of this has been observed in intercalated neurons of the amygdala [24].

Interestingly, such homeostatic regulation is also consistent with findings at the ultra-structural level. The physical size of a synapse, in particular the surface area of the postsynaptic density (PSD), is commonly used as a proxy for a synapse's efficacy. Bourne and Harris have reported coordinated changes in PSD surface areas of dendritic spines in the hippocampus after LTP induction such that increases in the PSD surface areas of synapses or the creation of new synapses seem to be balanced by decreases of PSD surface areas of other synapses or their complete elimination such that the total amount of PSD surface area stays approximately constant [3].

A proxy of synaptic efficacy that is more precise than PSD surface area is the number of α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptors (AMPARs) inside the PSD, however. This number is modulated during various forms of plasticity. Therefore, a full understanding of synaptic plasticity requires a careful description of the mechanisms that regulate AMPAR numbers in synapses.

Here we show how the behavior of keeping the sum of synaptic efficacies approximately constant on short time scales naturally arises from a generic model in wich individual synapses compete for a limited supply of synaptic building blocks such as AMPARs or other protein complexes that are necessary to stabilize AMPARs inside the PSD. We assume that these building blocks are distributed across the dendritic tree and enter and leave dendritic spines in a stochastic fashion. The model predicts that the redistribution of synaptic efficacies should act multiplicatively, as is often assumed in *ad hoc* normalization models. We also show that this model naturally gives rise to a homeostatic form of heterosynaptic plasticity, where synapses grow at the expense of other synapses. Overall, by considering the trafficking of synaptic building blocks inside a neuron's dendritic tree, the model offers a parsimonious and unified account of these fundamental plasticity mechanisms.

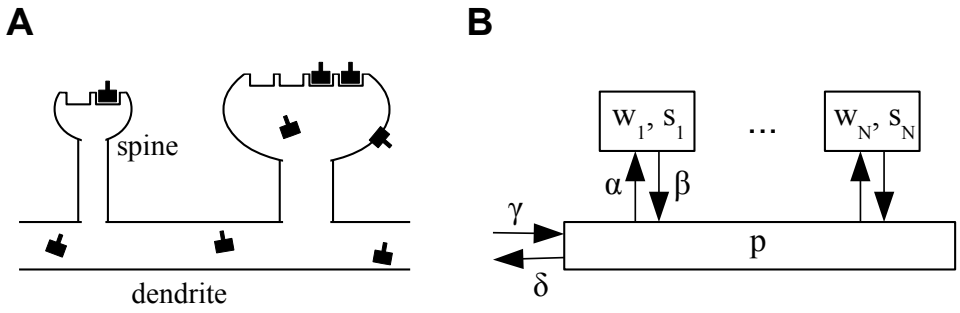

Figure 1. A. Sketch of the architecture of the model. Neurotransmitter receptors, e.g. AMPA receptors, are trafficked through the dendrite and bind to "slots" inside of dendritic spines. The efficacy of a synapse is assumed to be proportional to the number of receptors attached to its slots. **B.** Abstract description of the stochastic process indicating the rates at which receptors move in and out of slots in the synapses and the receptor pool in the dendrite. See text for details.

Methods

The architecture of the model is shown in Fig. 1. We consider a single neuron with $N \in \mathbb{N}$ synaptic inputs. Each synapse is characterized by two variables. First, each synapse $i \in 1, \ldots, N$ has a number of slots $s_i \in \mathbb{R}^{\geq 0}$ for neurotransmitter receptors. Second, at any time a certain number of slots w_i actually contain a receptor. $w_i \in \mathbb{R}^{\geq 0}$ determines the current weight or efficacy of a synapse. We assume that the PSD cannot hold more functional receptors than there are slots, i.e., $w_i \leq s_i$. Next to receptors in the synapses, the neuron maintains a pool of receptors in its dendritic tree. The size of this pool is denoted $p \in \mathbb{R}^{\geq 0}$. Note that for mathematical convenience we here consider the s_i, w_i and p to be real numbers that can take non-integer values.

Receptors can transition from the pool to empty slots in a synapse or detach from such a slot and return into the pool with rates $\alpha \in \mathbb{R}^{>0}$ and $\beta \in \mathbb{R}^{>0}$, respectively. Receptors in the pool are removed with a rate $\delta \in \mathbb{R}^{>0}$. To counteract this loss, new receptors are added at a rate $\gamma \in \mathbb{R}^{>0}$ and injected into the pool. In the limit of large receptor numbers, the dynamics of the system can be described by the following system of coupled ordinary nonlinear differential equations:

$$\dot{w}_i = -\beta w_i + \alpha p (s_i - w_i) \,, \; i = 1, \ldots, N \tag{1}$$

$$\dot{p} = -\delta p + \gamma + \sum_i \beta w_i - \sum_i \alpha p (s_i - w_i) \,. \tag{2}$$

In the first equation, $-\beta w_i$ describes the return of receptors from the synapse into the pool, and $\alpha p(s_i - w_i)$ describes the binding of receptors from the pool to empty slots in the synapse. In the second equation, $-\delta p$ describes the deletion of receptors from the pool, γ represents the gain of new receptors, $\sum_i \beta w_i$ describes the return of receptors from the synapses into the pool, and finally $-\sum_i \alpha p(s_i - w_i)$ describes the loss of receptors from the pool which bind to free slots in the synapses. Together, this is a system of $N + 1$ coupled ordinary nonlinear differential equations, which we will study analytically and numerically below. The simulation software was

written in Python and is available upon request.

The model can be interpreted in different ways. Its generic interpretation is that the "receptors" of the model are AMPA receptor (AMPAR) complexes composed of AMPARs and transmembrane AMPAR regulatory proteins (TARPs) such as stargazin. The "slots" are postsynaptic density structures comprising membrane-associated guanylate kinase (MAGUK) proteins such as PSD-95 attached to the postsynaptic membrane, which stabilize AMPA receptor complexes in the postsynaptic density (PSD) [10, 27, 30]. In this generic interpretation of the model, the pool of receptors is the set of AMPAR complexes that diffuse in the plasma membrane and that are captured by the slots. Addition of receptors to the pool then subsumes (some or all of) the processes that assemble AMPAR complexes and prepare them for the insertion into slots: assembly of the receptors from the component subunits, trafficking, attachment of TARPs, exocytosis, and potentially phosphorylation. Removal from the pool similarly subsumes the set of reverse processes. Several variations of this generic interpretation are possible depending on what exactly we would like to associate with the "receptors" in the model: AMPARs, AMPAR+TARP complexes, AMPAR+TARP complexes that have already been exocytosed, phosphorylated, etc. Essentially, our model is a two step model (production and insertion), but we leave it open for interpretation, what steps in the full chain of events are considered the "production" (subsumed in rate γ) and which steps are considered the "insertion" (subsumed in rate α).

Evidently, receptor slots themselves must also be stabilized inside the PSD somehow. A second, maybe somewhat counter-intuitive, interpretation of the model is therefore that it describes the binding and unbinding of receptor slots to what one might consider a *slot for a receptor slot* or simply *slot-for-a-slot*. In this interpretation of the model, the "receptors" in the description above are actually the PSD-95 slot proteins and the "slots" are slots-for-a-slot to which the PSD-95 proteins can attach. The model then describes the trafficking of PSD-95 into and out of the PSD, assuming that available AMPAR complexes are quickly redistributed among PSD-95 slots (compared to the time scale of addition and removal of these PSD-95 slots to the PSD). This interpretation may be particularly useful if the supply of PSD-95 is the limiting factor in determining the number of functional AMPARs bound inside the PSD [27]. We leave open the question what exactly the slots-for-a-slot might be. It is clear however, that PSD-95 molecules can form stable lattices inside the PSD such that PSD-95 proteins could act as slots for other PSD-95 proteins. Furthermore, they can bind both to actin filaments and transmembrane proteins such as neuroligins, which themselves bind to neurolexins in the presynaptic membrane to ensure the proper alignment of pre- and postsynaptic structures [26, 32]. Which interpretation of the model is most useful will depend on the experimental situation and the relative abundances and binding rates of the various players (AMPARs, PSD-95 proteins, TARPs, ...). A more complete model would explicitly describe all of the various interactions between these molecular players, of course, but would be more complex and require knowledge of the various reaction rates and is left for future work.

The analysis of the model presented in the following does not depend on which interpretation is chosen. The only additional assumption we will make is that there is a separation of time scales between the fast trafficking of the "receptors" into and out of the "slots" and the slow addition and removal of receptors to the pool. Our main results only depend on this qualitative feature of the model. For the first, generic, interpretation of the model the assumption of a separation of time scales appears justified. If we interpret the receptor pool of the model to comprise AMPARs

that have been exocytosed and diffuse in the cell membrane, then the half-life of an AMPAR in the pool is of the order of 10 minutes suggesting $\delta^{-1} = 10\,\mathrm{min}/\ln 2 \approx 14\,\mathrm{min}$ [11, 12]. In contrast, the time an AMPAR stays inside the PSD, which we interpret as the time the AMPAR is bound to a slot, appears to be of the order of maybe 30 seconds [8], suggesting $\beta^{-1} = 30\,\mathrm{s}/\ln 2 \approx 43\,\mathrm{s}$. Regarding the second, slots-for-a-slot, interpretation of the model, we note that the half-life of PSD-95 residing inside the synapse is of the order of 5 hours [29], implying $\beta^{-1} \approx 5\,\mathrm{h}/\ln 2 \approx 7\,\mathrm{h}$. In contrast, the global half-life of PSD-95 has been estimated to be 3.67 days [5], implying $\delta^{-1} = 3.67\,\mathrm{d}/\ln 2 \approx 5.30\,\mathrm{d}$. In either case, the assumption of a separation of time scales appears justified.

Results

Steady-State Analysis of the Basic Model: Multiplicative Scaling

We begin our analysis by finding the stationary solution of the system of coupled differential equations defined by (1) and (2). First, it is convenient to introduce the total number of synaptic slots $S \equiv \sum_i s_i$ and the total number of docked receptors or total synaptic weight $W \equiv \sum_i w_i$ and note that its time derivative is $\dot{W} = \sum_i \dot{w}_i$. This allows us to rewrite (2) as:

$$\dot{p} = -\delta p + \gamma + \beta W - \alpha p(S - W) \,. \tag{3}$$

To find the fixed point solution p^∞, w_i^∞ with $W^\infty = \sum_i w_i^\infty$, we set the time derivatives to zero, i.e., we require $\dot{w}_i = 0\ \forall i$ and $\dot{p} = 0$ above. Inserting the first condition into (1) and summing over i yields:

$$0 = -\beta W^\infty + \alpha p^\infty (S - W^\infty) \,. \tag{4}$$

Similarly, setting $\dot{p} = 0$ in (3) gives:

$$0 = -\delta p^\infty + \gamma + \beta W^\infty - \alpha p^\infty (S - W^\infty) \,. \tag{5}$$

Adding (4) to (5) then gives the solution for p^∞:

$$p^\infty = \frac{\gamma}{\delta} \,, \tag{6}$$

The simple and intuitive result is therefore that the total number of receptors in the pool in the steady state is given by the ratio of the production rate γ and the removal rate δ. Specifically, the presence of many receptors in the pool requires $\gamma \gg \delta$.

We now solve for the steady state solutions w_i^∞ of the w_i by again setting $\dot{w}_i = 0$ in (1) and using (6) to give:

$$w_i^\infty = \frac{1}{1 + \frac{\beta\delta}{\alpha\gamma}} s_i \equiv F s_i \,. \tag{7}$$

Importantly, we find $w_i^\infty \propto s_i$, i.e. in the steady state the weights of synapses are proportional to the numbers of slots they have. The constant of proportionality is a *filling fraction* and we denote it by F. Interestingly, the filling fraction F is independent of the number of receptor slots. Figure 2A plots F as a function of the ratio of the four rate constants $(\beta\delta)/(\alpha\gamma)$. Note that a filling fraction close to one requires $\beta\delta \ll \alpha\gamma$.

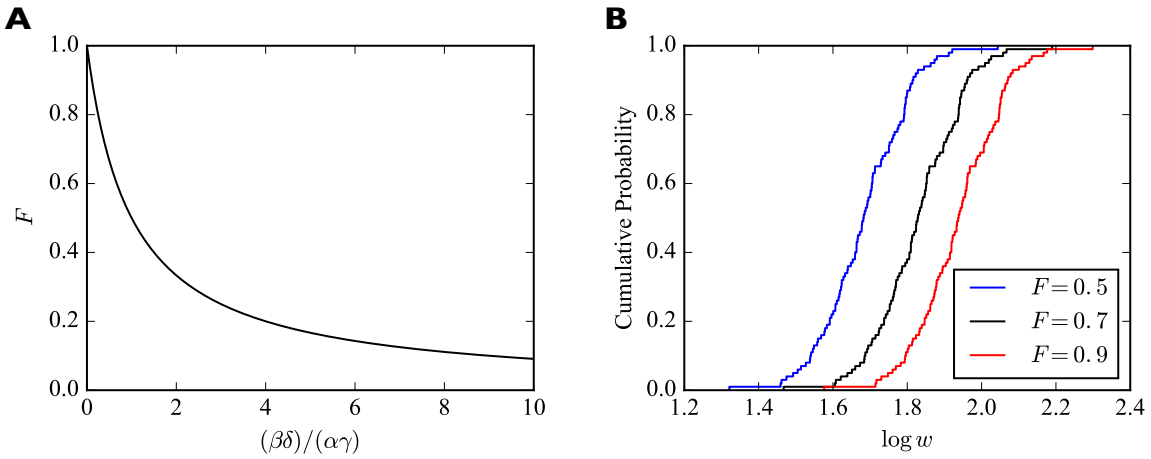

Figure 2. A. Filling fraction F as a function of the ratio of the rate constants $(\beta\delta)/(\alpha\gamma)$. **B.** Example empirical cumulative distribution functions (CDFs) of the numbers of receptors bound in individual synapses for fixed numbers of slots drawn from a lognormal distribution and three different filling fractions F. The simulated piece of dendrite has 100 synapses and 100 receptor slots per synapse on average.

Summing (7) over i reveals that $W^\infty = FS$, so we can also write:

$$w_i^\infty = \frac{s_i}{S} W^\infty \,, \tag{8}$$

where s_i/S is the relative contribution of synapse i to the total number of slots. Note that if the filling fraction changes, say, due to an increase in receptor production or a change in any of the other parameters, the *relative* strength of two synapses in the steady state is unaffected:

$$\frac{w_i^\infty}{w_j^\infty} = \frac{s_i}{s_j} = \text{const.} \tag{9}$$

Therefore, all synaptic efficacies will be scaled multiplicatively by the same factor.

To illustrate the effect of multiplicative scaling of synaptic efficacies, we consider a piece of dendrite with $N = 100$ afferent synapses. The number of receptor slots s_i in these synapses are drawn from a lognormal distribution with mean 1.0 and standard deviation 0.2 and subsequently scaled such that there are 100 slots per synapse on average. We consider three different filling fractions $F \in \{0.5, 0.7, 0.9\}$. The empirical cumulative distribution functions (CDFs) of the common (decadic) logarithms of the w_i are shown in Fig. 2A. The horizontal shifting of the empirical CDFs illustrates the multiplicative scaling of the individual synaptic efficacies.

The total number of receptors in the system in the steady state R^∞ is given by the sum of the number of receptors in the pool and the number of receptors attached to slots. Combining the above results, we find:

$$R^\infty = p^\infty + W^\infty = p^\infty + FS = \frac{\gamma}{\delta} + \frac{1}{1 + \frac{\beta\delta}{\alpha\gamma}} S \,. \tag{10}$$

In particular, the total number of receptors in the steady state depends on the total number of slots. The number of receptor slots also influences which fraction of receptors are in the pool vs. bound to slots. Specifically, we find for the steady state fraction of receptors residing in the pool:

$$\frac{p^\infty}{R^\infty} = \frac{p^\infty}{p^\infty + W^\infty} = \frac{\alpha\gamma + \beta\delta}{\alpha\gamma + \beta\delta \ + \ \alpha\delta S} \ . \tag{11}$$

Analogously, for the steady state fraction of receptors bound to slots we find:

$$\frac{W^\infty}{R^\infty} = 1 - \frac{p^\infty}{R^\infty} = \frac{\alpha\delta S}{\alpha\delta S \ + \ \alpha\gamma + \beta\delta} \ . \tag{12}$$

If we view these expressions as functions of S, assuming all other parameters as being fixed, then the intuitive result is that for small numbers of slots (S approaching zero), almost all receptors will remain in the pool. Conversely, as the number of slots grows bigger and the number of receptors populating these slots, FS, increases, the percentage of receptors in the pool declines.

The filling fraction F depends on all four rate parameters $\alpha, \beta, \gamma, \delta$. Due to lack of empirical data we treat γ as a free parameter that we can use to set the steady state pool size p^∞ via $\gamma = \delta p^\infty$. Similarly, we treat α as a free parameter that we can use to control the filling fraction. For fixed positive values of β, γ, δ, the filling fraction will approach zero when α approaches zero and it will approach one if α grows very big such that $\alpha\gamma \gg \beta\delta$. To set a desired filling fraction, we can solve the implicit definition of F from (7) for α to obtain:

$$\alpha = \frac{\beta\delta}{\gamma}\frac{F}{1-F} = \frac{\beta}{p^\infty}\frac{F}{1-F} \ . \tag{13}$$

The above results fully describe the system after it had sufficient time to reach its equilibrium. On a shorter time scale, however, the system may transiently assume different quasi-stationary states, because receptor addition and removal are slow compared to receptor binding and unbinding to and from slots. In the following, we consider the short-term behavior of the model on time scales where the total number of receptors is approximately constant. This will allows us to reveal, among other things, a transient form of heterosynaptic plasticity.

Short-Term Behavior for Approximately Constant Number of Receptors

In the following, we assume that the processes of receptor addition and removal to and from the pool are slow compared to the attaching and detaching of receptors to and from slots. For instance, the time that an AMPAR remains in the cell membrane is of the order of ten minutes while the time it resides inside the PSD is maybe of the order of half a minute. A reasonable approximation on short times scales is therefore to neglect the production and removal terms in (2). In this case, the total number of receptors $R \equiv W + p$ is constant, as can be seen by removing the $-\delta p$ and $+\gamma$ terms from (2), and adding (1), summed over all i, which gives $\dot{p} + \dot{W} = \dot{R} = 0$. We again look for the steady state solution by setting the time derivatives of w_i and p to zero and summing over i. This leads to the following quadratic equation for W^*, the steady state number of bound receptors in the short-term approximation (which must not be confused with the long-term steady state solution W^∞ of the full system):

$$W^{*2} - W^*\left(S + R + \frac{\beta}{\alpha}\right) + RS = 0 \ . \tag{14}$$

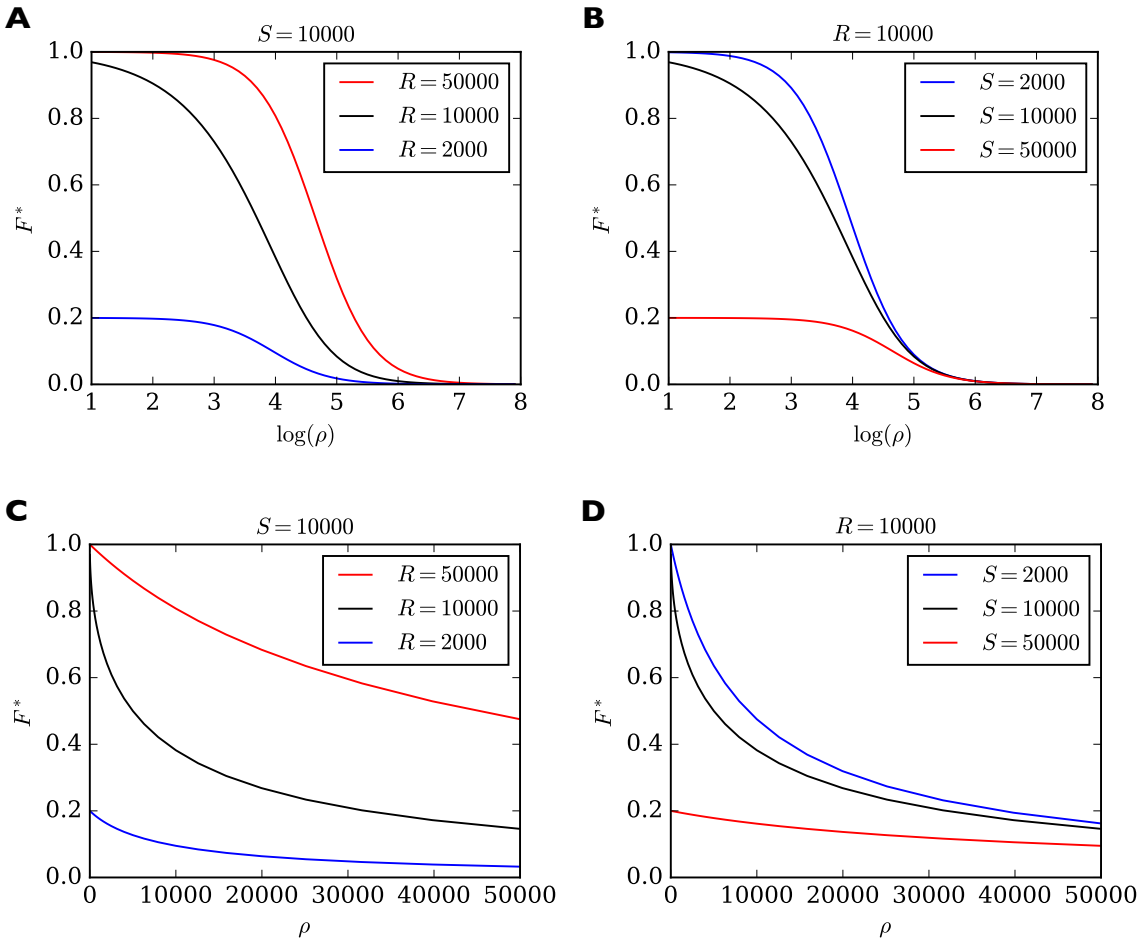

Figure 3. Filling fraction F^* in the short-term approximation of constant receptor number as a function of the ratio of transition rates $\rho = \beta/\alpha$ for different combinations of R and S. **A.** F^* for a fixed number of $S = 10\,000$ slots and three different receptor numbers as a function of $\log(\rho)$ **B.** For fixed number of $R = 10\,000$ receptors and three different numbers of slots. **C.,D.** Same as **A.,B.** but F^* is plotted as a function of ρ rather than $\log(\rho)$. Note that F^* reacts particularly sensitively to changes in ρ when ρ is small and when $R = S$ (black curves). In this regime, small changes to, say, the rate of detaching from slots β have a great influence on the filling fraction. In all cases, the shown solution F^* is only transient. Eventually the filling fraction will assume its steady state value F given by (7).

We introduce $\rho \equiv \beta/\alpha$ as the ratio of the rates through which receptors leave and enter the synaptic slots. Using this, the two solutions of (14) are given by:

$$W^*_{1,2} = \frac{1}{2}\left(S + R + \rho\right) \pm \sqrt{\frac{1}{4}\left(S + R + \rho\right)^2 - RS}\,. \tag{15}$$

The "+" solution is not biologically meaningful, since it leads to $W^* \geq S$ or $W^* \geq R$ (see

Appendix), so that the desired steady state solution of the short-term approximation is given by:

$$W^* \equiv W_2^* = \frac{1}{2}\left(S + R + \rho\right) - \sqrt{\frac{1}{4}\left(S + R + \rho\right)^2 - RS} \tag{16}$$

and the corresponding short-term steady-state filling fraction is $F^* = W^*/S$. In Fig. 3 we show the behavior of F^* as a function of ρ for different combinations of total number of slots S and total number of receptors R. F^* has an inverted sigmoidal shape as a function of $\log(\rho)$. For high values of ρ the filling fraction F^* always goes to zero. For ρ approaching zero, F^* achieves a maximum value which depends on whether there are fewer or more receptors than slots in the system. If there are more receptors than slots then F^* approaches one. If there are fewer receptors than slots then F^* approaches the ratio of receptors to slots in the system. In general, we find that the maximum short-term filling fraction for $\rho \to 0$ is given by $F^*_{\max} = \min\{1, R/S\}$. In particular, a high filling fraction can only be achieved if $R > S$, i.e., there must be more receptors than slots in the system. On the other hand, F^* is most sensitive to changes in ρ when $R = S$. This can be seen by the steep negative slope of the black curves in Fig. 3C,D for small values of ρ. Formally, we consider the partial derivative of the short-term filling fraction $F^* = W^*/S$ with respect to ρ. Using (16) we find:

$$\frac{\partial F^*}{\partial \rho} = \frac{1}{S}\frac{\partial W^*}{\partial \rho} = \frac{1}{S}\left[\frac{1}{2} - \frac{R + S + \rho}{4\sqrt{\frac{1}{4}(R + S + \rho)^2 - RS}}\right] . \tag{17}$$

As can be seen in Fig. 3C,D, the most extreme slope is obtained at $\rho = 0$. There the derivative simplifies to:

$$\left.\frac{\partial F^*}{\partial \rho}\right|_{\rho=0} = \frac{1}{2S}\left(1 - \frac{R + S}{R - S}\right) . \tag{18}$$

Importantly, for $R = S$ the derivative diverges, i.e., F^* reacts extremely sensitively to changes in ρ. For a very sensitive control of F^* over its full range it is therefore beneficial to have $R \approx S$. In the following, we study some of the consequences of the fast exchange of receptors between the pool and the synaptic slots.

Short-Term Response to Change in the Pool Size

To illustrate the fast redistribution of receptors, we consider a sudden change in the pool size. In our generic interpretation of the model, this corresponds to the sudden exocytosis or endocytosis of AMPARs. To study the effect of such a manipulation, we discretize the full dynamic equations using the Euler method and solve them numerically. For illustration, we consider a piece of dendrite with just three synapses with 40, 60, and 80 slots, whose pool size is changed abruptly (Fig. 4). Parameters are set to achieve a filling fraction of $F = 0.9$ and a steady state pool size $p^\infty = 100$. After 2 minutes, the number of receptors in the pool is either doubled (solid lines) or set to zero (dotted lines). In response, all synapses are rapidly scaled up or down multiplicatively. The new equilibrium is only transient, however. On a slower time scale the system returns to its starting point as the slow production and removal processes drive the system back to its steady state solution w_i^∞, p^∞.

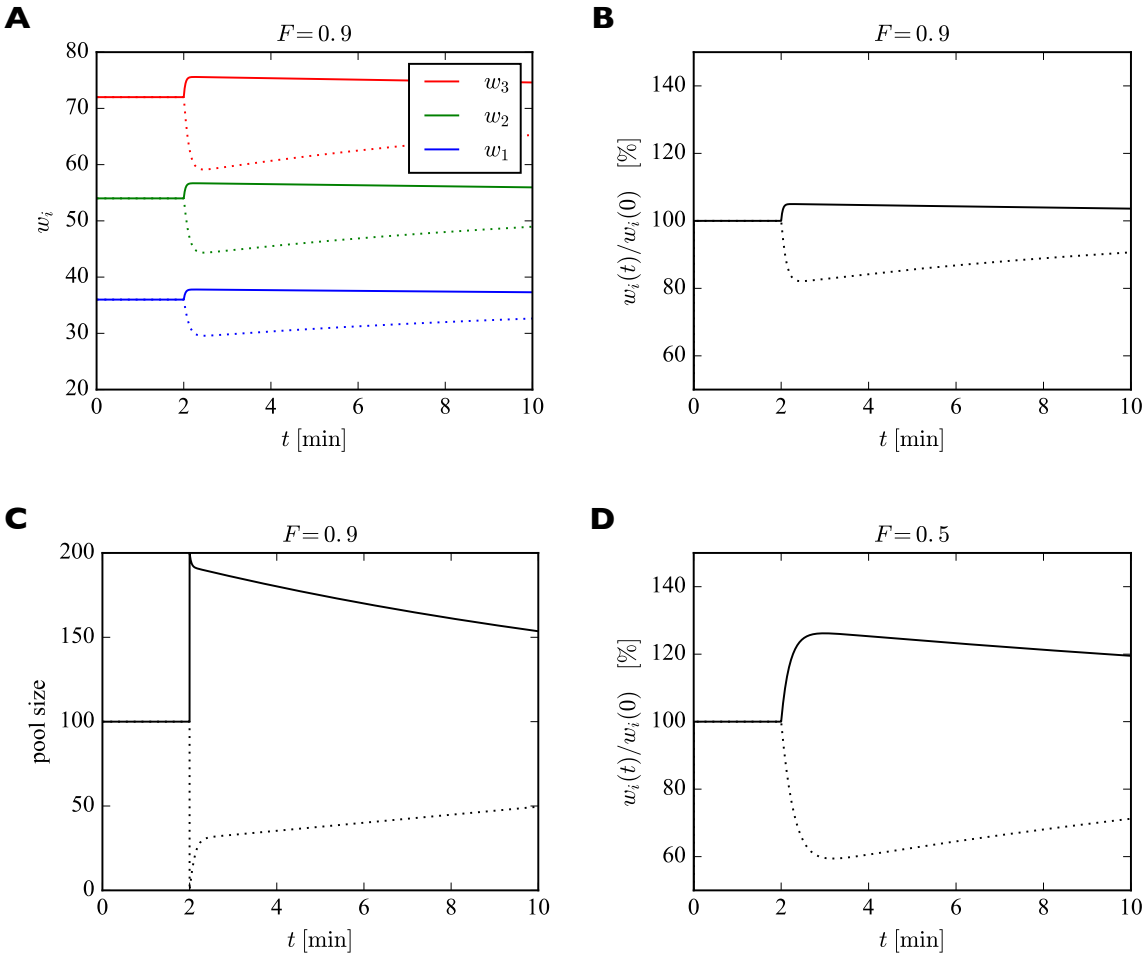

Figure 4. Effect of sudden change of the pool size p on synaptic efficacies. **A.** After 2 minutes, the pool size is either doubled (solid curves) or set to zero (dotted curves). In response, the synaptic efficacies are scaled multiplicatively as receptors are redistributed through the system. Doubling the receptor pool has a relatively weak effect in this example, as the system starts with a high filling fraction of 0.9, meaning that 90% of the slots are already filled at the beginning and there are few empty slots to which the additional receptors can bind. **B.** Same as **A.** but showing relative change in synaptic efficacies, which is identical for all synapses. **C.** Change in pool size. After the sudden increase or decrease in pool size at 2 minutes, there is first a rapid relaxation of the pool size followed by a much slower return towards the original value. **D.** Same as **B.** but for a filling fraction of 1/2. The smaller filling fraction leads to bigger relative changes of the synaptic efficacies. Parameters used were: $\beta = 1/43\ \mathrm{s}^{-1}$, $\delta = 1/14\ \mathrm{min}^{-1}$. The production rate γ was set to achieve a steady state pool size of $p^{\infty} = 100$ by setting $\gamma = 100 \times \delta$. α was chosen to achieve a filling fraction of $F = 0.9$ ($F = 0.5$ in **D.**) using (13).

The fast equilibration process to a transient steady state also naturally gives rise to a homeostatic form of heterosynaptic plasticity. When, e.g., the number of receptor slots in some synapses is quickly increased, then receptors are redistributed such that the efficacies of synapses with an increased number of receptor slots will grow, while the efficacies of other synapses will shrink, as we will discuss in the following.

Heterosynaptic Plasticity due to Competition for Receptors

The model gives rise to a form of heterosynaptic plasticity since all synapses are competing for a limited number of receptors inside the dendritic receptor pool. For illustration purposes we consider a neuron with four synaptic inputs (Fig. 5). At the beginning of the simulation, the number of slots in the four synapses are 20, 40, 60, and 80. We start the system in its steady state with a pool size of 20 and a filling fraction $F^\infty = 0.9$ leading to a total of $20 + 200 \times 0.9 = 200$ receptors in the system. After 2 minutes we instantaneously increase the number of slots in the first (blue) and third (red) synapse by 20%. Subsequently, the system settles into a new (transient) equilibrium (Fig. 5**A**). While w_1 and w_3 increase, the number of receptors in synapses 2 and 4 (and in the pool) slightly decrease, although their numbers of slots have not changed. This behavior corresponds to a form of heterosynaptic plasticity where synapses grow at the expense of other synapses and is due to the approximately constant number of receptors on a fast time scale. Note that the sum of synaptic efficacies is not perfectly constant in this example, however. The increase in synaptic efficacies w_1 and w_3 is bigger than the decrease of synaptic efficacies w_2 and w_4. This is because the new receptors entering synapses 1 and 3 are partly recruited from the receptor pool. The bigger the size of the pool, the stronger is this effect. Close to perfect balancing of synaptic weights would require $p \ll W$. Figure 5**B** shows the relative changes of efficacies of the synapses undergoing homosynaptic LTP (blue curve corresponding to w_1 and w_3 in **A**) vs. heterosynaptic LTD (green curve corresponding to w_2 and w_4 in **A**).

What determines the magnitude of the heterosynaptic change? We can calculate this analytically by using the above short-term approximation F^* for constant receptor number. Before plasticity induction, the synaptic efficacy of a synapse in equilibrium is given by $w_i = Fs_i = (s_i/S)W$. The induction of homosynaptic plasticity in other synapses changes the total number of available receptor slots and we denote the new number of slots S'. Shortly after homosynaptic plasticity induction the synaptic efficacies of a synapse that did not undergo homosynaptic plasticity will be approximately $w_i^* = F^* s_i = (s_i/S')W^*$ as receptors are redistributed through the system. Therefore, the relative heterosynaptic change of such a synapse is given by $(w_i^* - w_i)/w_i = (F^* - F)/F$ as long as the total number of receptors has not changed much.

In Fig. 5**C** we plot this relative change in synaptic efficacy due to heterosynaptic plasticity as a function of the total number of receptor slots following homosynaptic plasticity induction for different filling fractions. The pool size is assumed to be 10% of the total number of slots prior to plasticity induction. First, we can observe that reductions in the total number of slots due to homosynaptic LTD cause heterosynaptic LTP. Conversely, increases in the total number of slots due to homosynaptic LTP cause heterosynaptic LTD. Second, the amount of heterosynaptic plasticity depends on the filling fraction prior to plasticity induction. Specifically, a high filling fraction of 0.9 leads to weaker heterosynaptic LTP. As most slots are already filled prior to plasticity induction, the potential for undergoing heterosynaptic LTP is small, as there are few

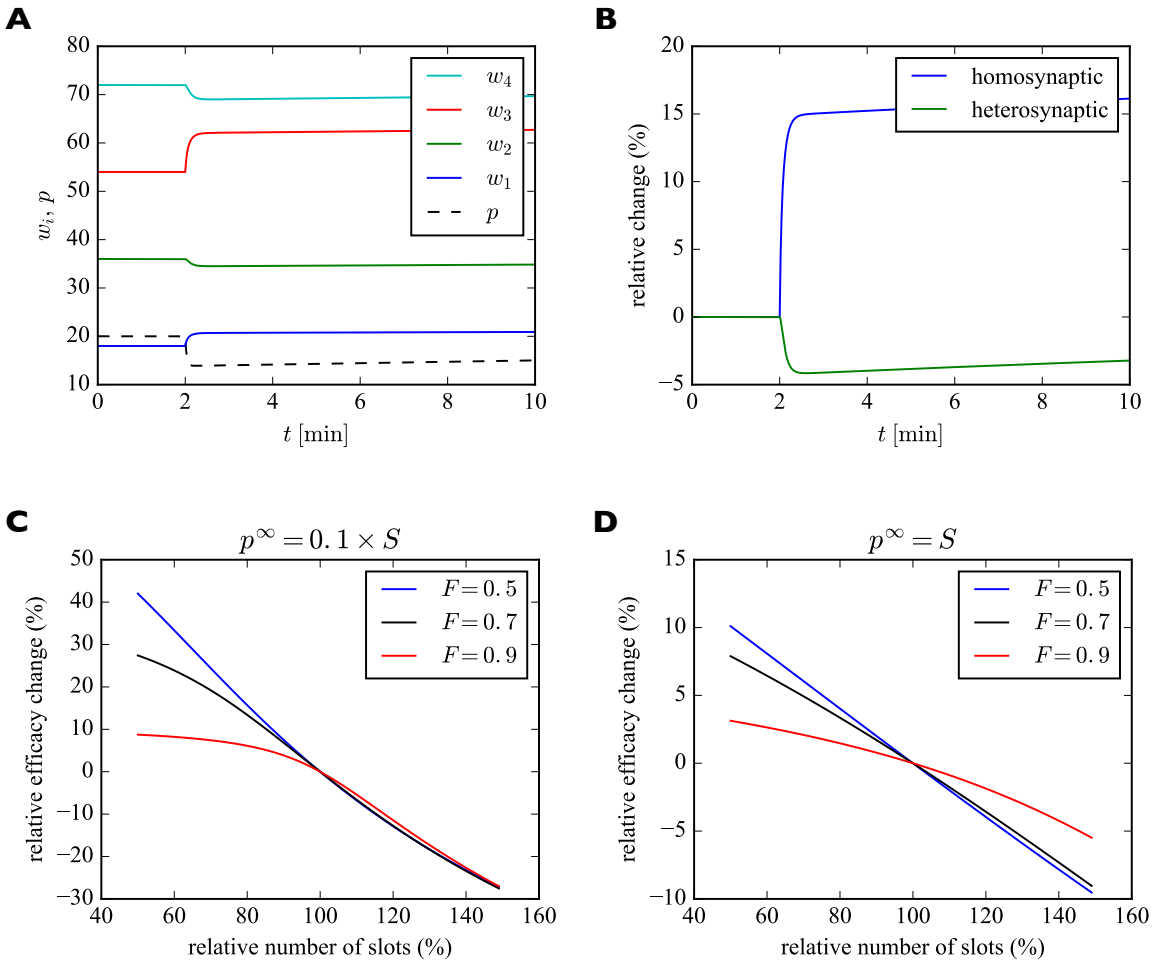

Figure 5. A. Illustration of heterosynaptic plasticity. After 2 minutes, the number of slots in synapses 1 and 3 is increased instantaneously. The system quickly reaches a new (transient) equilibrium, where the non-stimulated synapses 2 and 4 are slightly weakened. At the same time, the number of receptors in the pool is reduced. Parameters were: $\beta = 1/43\ \mathrm{s}^{-1}$, $\delta = 1/14\ \mathrm{min}^{-1}$. The production rate γ was set to achieve a steady state pool size of $p^\infty = 20$ by setting $\gamma = 20 \times \delta$. α was chosen to achieve a filling fraction of $F = 0.9$ using (13). **B.** Time course of relative changes in synaptic efficacies due to homosynaptic and heterosynaptic plasticity for the experiment from **A**. **C.** Approximate maximum relative change of synaptic efficacy due to heterosynaptic plasticity as a function of the number of receptor slots after homosynaptic plasticity induction for different filling fractions. **D.** Same as **C** but for a large receptor pool. See text for details.

empty slots available in the synapse.

Figure 5**D** shows the analogous solution for the case of a big receptor pool. Everything else is identical to Fig. 5**C**. The strongly filled pool greatly attenuates the heterosynaptic plasticity

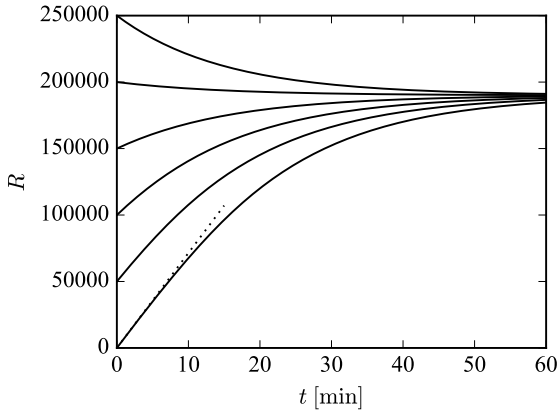

Figure 6. Illustration of long-term behavior under separation of time scales assumption. The neuron has a total of 100 000 receptor slots and is initialized with different receptor numbers. For low receptor numbers, the growth rate of R is approximately γ (compare dotted line). For a filling fraction close to its final steady-state value, R exponentially converges to its steady state of $\gamma/\delta + FS$ with time constant δ^{-1}. Parameters were: $\beta = 1/43\ \mathrm{s}^{-1}$, $\delta = 1/14\ \mathrm{min}^{-1}$. The production rate γ was set to achieve a steady state pool size $p^\infty = 100\,000$ by setting $\gamma = 100\,000 \times \delta$. α was set to achieve a final filling fraction of $F = 0.9$ using (13), so that the steady-state total number of receptors is $R^\infty = p^\infty + FS = 190\,000$.

effect (note the different scaling on the y-axes of **C** and **D**). When, e.g., new slots are added in this case, they will recruit receptors from the large receptor pool and the heterosynaptic effect on other synapses remains small. The large pool essentially functions as a buffer and somewhat shields synapses from heterosynaptic plasticity. Consistent with **C**, larger filling fractions again lead to less heterosynaptic plasticity.

Importantly, these effects are inherently transient. Over a sufficiently long time, the system will settle into a new (true) equilibrium, where every synapse has the same filling fraction F^∞ determined by the rate constants $\alpha, \beta, \gamma, \delta$ as described above. In the particular example of Fig. 5**A**,**B**, synapses 2 and 4 slowly return to their original efficacies, while synapses 1 and 3 remain permanently strengthened due to their increased number of slots. This effect might explain the often transient nature of heterosynaptic plasticity observed in experiments, e.g., [1].

Another mechanism for producing such a heterosynaptic effect is by changing the transition rates α and β in a synapse-specific fashion. For example, increasing α for some synapses will attract additional receptors to these synapses and lead to a heterosynaptic removal of receptors from the remaining synapses and the receptor pool (not shown).

This removal of receptors from the remaining synapses and the new equilibria seen in Figs. 4 and 5 are only transient. In the following, we will study the long-term behavior of the model on the time scale associated with receptor production and removal to determine how long it takes for the system to reach its stable fixed point given by w_i^∞ and p^∞.

Long-Term Behavior under Separation of Time Scales

To study the behavior of the system on the slow time scale of receptor production and removal we again make use of the separation of time scales argument. Specifically, we assume that the fast dynamics of receptor exchanges between the pool and the synapses quickly reaches its equilibrium before the total number of receptors can change much. The change in the total receptor number from (1) and (2) is approximated by:

$$\dot{R} = \dot{p} + \dot{W} = -\delta p + \gamma \approx -\delta p^* + \gamma \,, \tag{19}$$

where we have replaced the current pool size p with its steady state value $p^*(R) = R - W^*(R)$ for a constant number of receptors in the system. Using $F^*(R) \equiv W^*(R)/S$ we arrive at:

$$\dot{R} = \gamma + \delta F^*(R) S - \delta R \,. \tag{20}$$

For small numbers of receptors in the system, i.e. R close to zero, the steady state filling fraction $F^*(R)$ will be close to zero so that $\dot{R} \approx \gamma$. In contrast, for high numbers of receptors and the filling fraction close to its long-term steady-state value F we find:

$$\frac{1}{\delta}\dot{R} \approx \frac{\gamma}{\delta} + FS - R \,, \tag{21}$$

indicating that R will exponentially approach its steady state value of $\gamma/\delta + FS$ with the time constant $1/\delta$. This behavior is illustrated in Fig. 6, where we plot the numerical solution of (20) for different initial numbers of receptors in the system. Note the approximately linear growth of R with slope γ for small values of R (dotted line).

Discussion

The detailed molecular mechanisms underlying different forms of synaptic plasticity are complex. Recent years have seen enormous progress in identifying many of the relevant molecules and signaling pathways. This rapid development is in stark contrast to the simplistic and often purely phenomenological descriptions of synaptic plasticity used in most neural network models. While highly simplified mathematical models have been essential for relating synaptic plasticity "rules" to learning processes at the network level, a full understanding of synaptic plasticity requires the development of more elaborate models that do justice to the complexities of synaptic plasticity at the molecular level. Here we have taken a step in this direction. Our effort focuses on postsynaptic expression of plasticity and can be seen as complementing previous work on the mechanisms of presynaptic forms of plasticity [33].

Hebbian learning tends to lead to runaway growth of synaptic efficacies if not counteracted by competitive or homeostatic mechanisms. To be effective, these compensatory mechanisms must act fast enough so they can catch up with changes induced by Hebbian learning [4, 37]. Prominent candidate mechanisms are synaptic normalization and heterosynaptic plasticity [19]. The idea is a very old one. Synapses on the dendritic tree compete for limited resource of synaptic building blocks such that when some synapses grow, they have to do so at the expense of other synapses [19, 20]. Here we have presented a concrete model of such a fast normalization of the

efficacies of a neuron's afferent synapses based on a competition for synaptic resources such as AMPA receptor molecules.

Our model makes several contributions. First, it formalizes the idea of a fast synaptic normalization based on a competition for dendritic resources in an abstract and analytically tractable model. Second, analysis of the model reveals that under the given assumptions, normalization should act multiplicatively, such that relative strengths of synapses are maintained. Multiplicative normalization rules have been used in neural network models for a long time but usually in an *ad hoc* fashion. Our model supports the idea that a fast multiplicative normalization may in fact be biologically plausible. Third, the model naturally gives rise to a homeostatic form of heterosynaptic plasticity where, on a fast time scale, synapses grow in efficacy at the expense of other synapses.

The model also makes several predictions. First, the model predicts that the short-term filling fraction reacts most sensitively to changes in the rate constants for receptors moving into and out of receptor slots if the total number of receptors approximately equals the total number of receptor slots. This high sensitivity could be exploited for a very effective control of the filling fraction at the level of individual synapses, if these rate constants are controlled in a synapse-specific fashion (more on this below in the Section "Modeling early and late LTP and LTD"). Second, the model predicts that the amount of heterosynaptic plasticity depends on the size of the pool of available receptors. When this pool is big, heterosynaptic plasticity is reduced. Third, the model predicts that the amount of heterosynaptic plasticity depends on the current filling fraction of receptor slots. When, e.g., most receptor slots are already occupied due to a high filling fraction, then heterosynaptic LTP is reduced.

Despite these contributions, the model is only a first step towards a full mechanistic account of the molecular mechanisms underlying these different forms of synaptic plasticity. In the following we discuss a number of possible extensions to the model and make some additional predictions.

Stochastic version of the model. The present model was formulated as a system of nonlinear differential equations. This deterministic scheme is appropriate in the limit of large numbers of receptors, but it does not describe fluctuations of receptor numbers due to the stochastic attaching, detaching, addition, and removal of individual receptors. Such fluctuations become essential when receptor numbers are low, i.e., in small synapses. To better understand the effects of fluctuations, a stochastic version of the model is needed [6]. Such a model could also incorporate the effects of protein crowding in the PSD [15].

Dendritic morphology and local production. We have assumed that the transition rates α and β are identical for all synapses. This assumption is essential for the multiplicative behavior of the model and it may be warranted if the concentration of "free" receptors in the pool is more or less uniform across the dendritic tree. Otherwise one would expect, e.g., a higher effective α in regions of the dendritic tree with a high concentration of free receptors. Properly distributing synaptic building blocks across the dendritic tree is a formidable task [35]. Specifically, if receptors were produced at a single site corresponding to the cell nucleus and spreading from this point source according to a slow diffusion process then one would expect high concentration of

receptors close to the nucleus and low concentration far away from it. This would, all else being equal, lead to high α for synapses close to the cell nucleus and low α for synapses far away from it. Local production of synaptic building blocks such as AMPARs may lead to a more uniform distribution of these building blocks across the dendritic tree and therefore be essential for the multiplicative scaling behavior observed in the model and in biological experiments [34]. We therefore predict that local production of synaptic building blocks contributes to their uniform distribution across the dendritic tree which in turn guarantees multiplicative scaling behavior and the maintenance of relative strengths of synapses.

In this context it is also interesting to note that at least one form of heterosynaptic plasticity tends to be induced locally [7, 16, 18], i.e., at neighboring synapses. Such local action is readily expected in our model if competition arises for synaptic building blocks in a *local* pool, say, a dendritic branch, with comparatively slow trafficking of building blocks between multiple such local pools.

Balancing slots and receptors. For efficient functioning of the neuron the number of slots and the number of receptors need to be coupled. Having orders of magnitude more slots than receptors or vice versa would waste precious resources as the total synaptic weight is necessarily limited by the smaller of the two quantities. However, for a neuron to directly estimate its total number of slots may be difficult. An elegant solution for achieving this in a distributed fashion would be that the neuron regulates its receptor production based on the concentration of receptors in the pool, because if slots are added or removed the concentration of receptors in the pool will decrease or increase, respectively (compare Fig. 5). If, e.g., during development the neuron is growing new dendritic branches and forming new synapses on these then the (local) receptor concentration will be reduced. Increasing (local) production to maintain a constant concentration of receptors in the pool would counteract this effect and could ensure maintenance of a constant ratio between slots and receptors. We therefore predict the existence of a signaling pathway that regulates production of receptors based on their abundance in the receptor pool.

Slot production and removal. The above example raises the more general question of how slots are added and removed. Apart from the increase in the number of slots of single synapses in our illustration of heterosynaptic plasticity, we have treated the number of slots as fixed. In future work, it will be interesting to consider changes to slot numbers in more detail. Obviously, the building blocks of these "slots" also have to be produced, transported, and inserted into synapses, which could be based on similar mechanisms as we have postulated for receptors. Furthermore, slots may also be degraded and have to be replaced. In fact, the alternative interpretation of our model discussed in the beginning already describes how PSD-95 slots are produced (or degraded) and bind to (or detach from) slots for these receptor slots ("slots-for-a-slot"). Future work should aim for a model that more fully describes the interactions of AMPARs (and other types of receptors), various TARPs such as stargazin, MAGUK proteins such as PSD-95, and neuroligins as well as their production and trafficking.

Modeling early and late LTP and LTD. Along these lines, it will also be interesting to consider the mechanisms underlying different stages of LTP and LTD. A promising approach

would be the following. We have considered the rate constants α and β for receptors attaching to and detaching from slots to be constant and identical for all synapses. We can relax this condition to model the situation that due to the application of a plasticity induction protocol a synapse has been "tagged" for undergoing LTP or LTD as proposed by the *synaptic tagging and capture hypothesis* [9]. A simple way of formalizing this idea for, say, LTP in the model is the following: Once a synapse has been tagged for LTP the rate constant α for that synapse is transiently increased (and/or β may be decreased). This will quickly lead to a bigger filling fraction for this particular synapse, which is one mechanism of so-called early-LTP (next to receptor phosphorylation). In addition, we propose to add a process that slowly attracts and integrates new slots into the PSD, which is a mechanism of so-called late-LTP. When the tag is removed, α will return to its normal value and slot insertion will stop. The synapse will then return to its normal filling fraction but now with an increased number of slots. As a result, the synaptic efficacy has been permanently increased. When many synapses are tagged for LTP simultaneously, then a competition for new slots will ensue, analogous to the competition for receptors in the present model causing heterosynaptic plasticity.

Modeling slow homeostatic synaptic scaling. The model can also be extended to capture slow homeostatic synaptic scaling processes [13, 34]. In the simplest case, a sensor for the average neural activity of the neuron would drive the production of receptors and/or slots in a homeostatic fashion, such that if, e.g., the average neural activity falls below a target level or range, then receptor and/or slot production are increased to drive up excitatory synaptic efficacies. Such a model would naturally explain the multiplicative behavior of homeostatic synaptic scaling [34]. Obviously, the activity sensor could also sense the average activity in a local neighborhood through a diffusive mechanism [31]. Furthermore, instead of homeostatically regulating firing rates, the amount of afferent drive to the neuron or to the local population could be controlled [25], or even other measures of neural and synaptic activity could be used. Finally, all these ideas are not mutually exclusive. It seems likely that neurons control both their firing rate distributions and their amounts of excitatory and inhibitory afferent drive through a combination of different intrinsic and synaptic plasticity mechanisms.

Receptor subunit composition. Finally, not all AMPARs are created equal. Depending on the composition of subunits, AMPARs have distinct properties in terms of, e.g., calcium permeability and trafficking properties (see [12] for a recent review). A more complete model should incorporate the diversity of AMPARs (or even other receptor types) and their properties.

Conclusion. In conclusion, our model offers a parsimonious explanation for a homeostatic form of heterosynaptic plasticity and fast synaptic normalization, which it predicts to be multiplicative. It therefore supports the use of such rules in neural network models. Despite its simplicity, we hope that the model will provide a useful framework for future experimental and theoretical studies on the detailed molecular mechanisms of Hebbian and homeostatic plasticity across different time scales.

Acknowledgments

Jochen Triesch was supported by the Quandt foundation. Anne-Sophie Hafner is a postdoctoral fellow in Erin Schumans department at the Max Planck Institute for Brain Research supported by fundings from an EMBO fellowship ALTF 1095-2015 and a grant from the Alexander von Humboldt Foundation. We thank Christian Tetzlaff and Erin Schuman for helpful discussions.

Appendix

We show that the "+" solution from (15) is not biologically meaningful. To see this, first note that $W_1 \leq W_2$. Furthermore, any meaningful solution W must fulfill $W \leq R$ and $W \leq S$, i.e., the number of receptors bound to slots cannot be bigger than the total number of receptors or the total number of slots. If the smaller solution W_1 does not meet both criteria, then the larger W_2 cannot meet them either. So we assume in the following that W_1 meets both these criteria so that $W_1 \leq \min\{R, S\}$. Our argument uses Vieta's formulas for the quadratic equation (5):

$$W_1 + W_2 = R + S + \rho \quad \text{and} \quad W_1 W_2 = RS \; .$$

Using the second formula we can write:

$$RS = W_1 W_2 \leq \min\{R, S\} W_2 \; ,$$

from which follows that:

$$W_2 \geq \frac{RS}{\min\{R, S\}} \; .$$

In the case that $R > S$, this leads to $W_2 \geq R$. The only biologically meaningful solution to this is the equality $W_2 = R$. This is the extreme case where all receptors are bound in slots and no receptors remain in the pool. With Vieta's second formula we see that in this case $W_1 = S$. Plugging both results into Vieta's first formula, we see that this solution requires $\rho = 0$, which in turn requires $\beta = 0$. In this case, no receptors would ever leave synapses.

The case $R < S$ leads to $W_2 \geq S$. The only biologically meaningful solution to this is the equality $W_2 = S$. This is the extreme case where all slots are filled with receptors. Using Vieta's formulas again leads to the uninteresting requirement $\rho = 0$ for this solution.

Finally, the case $R = S$ leads to $R = S = W_1 = W_2$ and also requires $\rho = 0$. In summary, the "+" solution in (12) only admits the extreme solutions $W = S$ or $W = R$ requiring $\rho = 0$ (and therefore $\beta = 0$), which are not biologically meaningful.

References

1. WC Abraham and Graham V Goddard. Asymmetric relationships between homosynaptic long-term potentiation and heterosynaptic long-term depression. *Nature*, 305(5936):717–719, 1983.

2. Craig H Bailey, Maurizio Giustetto, Yan-You Huang, Robert D Hawkins, and Eric R Kandel. Is heterosynaptic modulation essential for stabilizing hebbian plasiticity and memory. *Nature Reviews Neuroscience*, 1(1):11–20, 2000.

3. Jennifer N Bourne and Kristen M Harris. Coordination of size and number of excitatory and inhibitory synapses results in a balanced structural plasticity along mature hippocampal CA1 dendrites during LTP. *Hippocampus*, 21(4):354–373, 2011.

4. Marina Chistiakova, Nicholas M Bannon, Jen-Yung Chen, Maxim Bazhenov, and Maxim Volgushev. Homeostatic role of heterosynaptic plasticity: models and experiments. *Frontiers in computational neuroscience*, 9:89, 2015.

5. Laurie D Cohen, Rina Zuchman, Oksana Sorokina, Anke Müller, Daniela C Dieterich, J Douglas Armstrong, Tamar Ziv, and Noam E Ziv. Metabolic turnover of synaptic proteins: kinetics, interdependencies and implications for synaptic maintenance. *PloS one*, 8(5):e63191, 2013.

6. Katalin Czöndör, Magali Mondin, Mikael Garcia, Martin Heine, Renato Frischknecht, Daniel Choquet, Jean-Baptiste Sibarita, and Olivier R Thoumine. Unified quantitative model of AMPA receptor trafficking at synapses. *Proceedings of the National Academy of Sciences*, 109(9):3522–3527, 2012.

7. Mathias De Roo, Paul Klauser, and Dominique Muller. LTP promotes a selective long-term stabilization and clustering of dendritic spines. *PLoS Biol*, 6(9):e219, 2008.

8. Michael D Ehlers, Martin Heine, Laurent Groc, Ming-Chia Lee, and Daniel Choquet. Diffusional trapping of glur1 ampa receptors by input-specific synaptic activity. *Neuron*, 54(3):447–460, 2007.

9. Uwe Frey and Richard GM Morris. Synaptic tagging and long-term potentiation. *Nature*, 385(6616):533, 1997.

10. Anne-Sophie Hafner, Andrew C Penn, Dolors Grillo-Bosch, Natacha Retailleau, Christel Poujol, Amandine Philippat, Françoise Coussen, Matthieu Sainlos, Patricio Opazo, and Daniel Choquet. Lengthening of the stargazin cytoplasmic tail increases synaptic transmission by promoting interaction to deeper domains of PSD-95. *Neuron*, 86(2):475–489, 2015.

11. Jeremy M Henley and Kevin A Wilkinson. AMPA receptor trafficking and the mechanisms underlying synaptic plasticity and cognitive aging. *Dialogues Clin Neurosci*, 15(1):11–27, 2013.

12. Jeremy M Henley and Kevin A Wilkinson. Synaptic AMPA receptor composition in development, plasticity and disease. *Nature Reviews Neuroscience*, 17(6):337–350, 2016.

13. Keiji Ibata, Qian Sun, and Gina G Turrigiano. Rapid synaptic scaling induced by changes in postsynaptic firing. *Neuron*, 57(6):819–826, 2008.

14. Andreea Lazar, Gordon Pipa, and Jochen Triesch. SORN: a self-organizing recurrent neural network. *Frontiers in computational neuroscience*, 3:23, 2009.

15. Tuo P Li, Yu Song, Harold D MacGillavry, Thomas A Blanpied, and Sridhar Raghavachari. Protein crowding within the postsynaptic density can impede the escape of membrane proteins. *Journal of Neuroscience*, 36(15):4276–4295, 2016.

16. Yinyun Li, Tomas Kulvicius, and Christian Tetzlaff. Induction and consolidation of calcium-based homo-and heterosynaptic potentiation and depression. *PloS one*, 11(8):e0161679, 2016.

17. Yonatan Loewenstein, Annerose Kuras, and Simon Rumpel. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *Journal of Neuroscience*, 31(26):9481–9488, 2011.

18. Attila Losonczy, Judit K Makara, and Jeffrey C Magee. Compartmentalized dendritic plasticity and input feature storage in neurons. *Nature*, 452(7186):436–441, 2008.

19. Gary S Lynch, Thomas Dunwiddie, and Valentin Gribkoff. Heterosynaptic depression: a postsynaptic correlate of long-term potentiation. *Nature*, 266(5604):737–739, 1977.

20. Christoph von der Malsburg. Self-organization of orientation sensitive cells in the striate cortex. *Biological Cybernetics*, 14(2):85–100, 1973.

21. Kenneth D Miller and David JC MacKay. The role of constraints in hebbian learning. *Neural Computation*, 6(1):100–126, 1994.

22. Daniel Miner and Jochen Triesch. Plasticity-driven self-organization under topological constraints accounts for non-random features of cortical synaptic wiring. *PLoS Comput Biol*, 12(2):e1004759, 2016.

23. Erkki Oja. Simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273, 1982.

24. Sébastien Royer and Denis Paré. Conservation of total synaptic weight through balanced synaptic depression and potentiation. *Nature*, 422(6931):518–522, 2003.

25. Cristina Savin, Jochen Triesch, and Michael Meyer-Hermann. Epileptogenesis due to glia-mediated synaptic scaling. *Journal of The Royal Society Interface*, 6(37):655–668, 2009.

26. Peter Scheiffele, Jinhong Fan, Jenny Choih, Richard Fetter, and Tito Serafini. Neuroligin expressed in nonneuronal cells triggers presynaptic development in contacting axons. *Cell*, 101(6):657–669, 2000.

27. Eric Schnell, Max Sizemore, Siavash Karimzadegan, Lu Chen, David S Bredt, and Roger A Nicoll. Direct interactions between PSD-95 and stargazin control synaptic AMPA receptor number. *Proceedings of the National Academy of Sciences*, 99(21):13902–13907, 2002.

28. Sen Song, Per Jesper Sjöström, Markus Reigl, Sacha Nelson, and Dmitri B Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol*, 3(3):e68, 2005.

29. James F Sturgill, Pascal Steiner, Brian L Czervionke, and Bernardo L Sabatini. Distinct domains within PSD-95 mediate synaptic incorporation, stabilization, and activity-dependent trafficking. *Journal of Neuroscience*, 29(41):12845–12854, 2009.

30. Akio Sumioka, Dan Yan, and Susumu Tomita. TARP phosphorylation regulates synaptic AMPA receptors through lipid bilayers. *Neuron*, 66(5):755–767, 2010.

31. Yann Sweeney, Jeanette Hellgren Kotaleski, and Matthias H Hennig. A diffusive homeostatic signal maintains neural heterogeneity and responsiveness in cortical networks. *PLoS Comput Biol*, 11(7):e1004389, 2015.

32. Ai-Hui Tang, Haiwen Chen, Tuo P Li, Sarah R Metzbower, Harold D MacGillavry, and Thomas A Blanpied. A trans-synaptic nanocolumn aligns neurotransmitter release to receptors. *Nature*, 2016.

33. Misha Tsodyks, Klaus Pawelzik, and Henry Markram. Neural networks with dynamic synapses. *Neural computation*, 10(4):821–835, 1998.

34. Gina G Turrigiano, Kenneth R Leslie, Niraj S Desai, Lana C Rutherford, and Sacha B Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391(6670):892–896, 1998.

35. Alex H Williams, Cian O'Donnell, Terrence J Sejnowski, and Timothy O'Leary. Dendritic trafficking faces physiologically critical speed-precision tradeoffs. *eLife*, 5, 2016.

36. Zhihua Wu and Yoko Yamaguchi. Conserving total synaptic weight ensures one-trial sequence learning of place fields in the hippocampus. *Neural networks*, 19(5):547–563, 2006.

37. Friedemann Zenke, Guillaume Hennequin, and Wulfram Gerstner. Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS Comput Biol*, 9(11):e1003330, 2013.

38. Pengsheng Zheng, Christos Dimitrakakis, and Jochen Triesch. Network self-organization explains the statistics and dynamics of synaptic connection strengths in cortex. *PLoS Comput Biol*, 9(1):e1002848, 2013.