

Fixed Points of Competitive Threshold-Linear Networks

Carina Curto

ccurto@psu.edu

Jesse Geneson

geneson@gmail.com

*Department of Mathematics, The Pennsylvania State University,
University Park, PA 16802, U.S.A.*

Katherine Morrison

katherine.morrison@unco.edu

*School of Mathematical Sciences, University of Northern Colorado,
Greeley, CO 80639, U.S.A.*

Threshold-linear networks (TLNs) are models of neural networks that consist of simple, perceptron-like neurons and exhibit nonlinear dynamics determined by the network's connectivity. The fixed points of a TLN, including both stable and unstable equilibria, play a critical role in shaping its emergent dynamics. In this work, we provide two novel characterizations for the set of fixed points of a competitive TLN: the first is in terms of a simple sign condition, while the second relies on the concept of domination. We apply these results to a special family of TLNs, called combinatorial threshold-linear networks (CTLNs), whose connectivity matrices are defined from directed graphs. This leads us to prove a series of graph rules that enable one to determine fixed points of a CTLN by analyzing the underlying graph. In addition, we study larger networks composed of smaller building block subnetworks and prove several theorems relating the fixed points of the full network to those of its components. Our results provide the foundation for a kind of graphical calculus to infer features of the dynamics from a network's connectivity.

1 Introduction ---

Threshold-linear networks (TLNs) are commonly used models of recurrent networks consisting of simple, perceptron-like neurons with nonlinear interactions. The dynamics are given by the system of ordinary differential equations,

$$\frac{dx_i}{dt} = -x_i + \left[\sum_{j=1}^n W_{ij}x_j + b_i \right]_+, \quad i = 1, \dots, n, \quad (1.1)$$

where n is the number of neurons, $x_i(t)$ is the activity level (or “firing rate”) of the i th neuron, W_{ij} is the connection strength from neuron j to neuron i , and $[\cdot]_+ \stackrel{\text{def}}{=} \max\{\cdot, 0\}$ is the threshold nonlinearity. The external inputs $b_i \in \mathbb{R}$ may be heterogeneous but are assumed to be constant in time. We refer to a given choice of TLN as (W, b) . Competitive TLNs have the additional requirement that all interactions are effectively inhibitory, with matrix entries $W_{ij} \leq 0$.

TLNs have previously been studied through the lens of permitted and forbidden sets (Hahnloser, Sarpeshkar, Mahowald, Douglas, & Seung, 2000; Xie, Hahnloser, & Seung, 2002; Hahnloser, Seung, & Slotine, 2003; Curto, Degeratu, & Itskov, 2012, 2013; Curto & Morrison, 2016), though this work was largely restricted to the case of symmetric networks. TLNs that are both symmetric and competitive generically exhibit multistability, with activity that is guaranteed to converge to a stable fixed point irrespective of initial conditions (Hahnloser et al., 2003). This property motivated the use of these networks as models for associative memory encoding and retrieval, similar to the Hopfield model (Hopfield, 1982). Earlier results on the mathematical theory of TLNs thus focused primarily on stable fixed points of symmetric networks. The papers (Hahnloser et al., 2000; Xie et al., 2002; Hahnloser et al., 2003) gave characterizations and applications of permitted sets, which are subsets of neurons that have the capacity to support a stable fixed point of equation 1.1 for some (potentially unknown) external input vector b . These authors also found nice properties satisfied by the full collection of permitted sets of a symmetric TLN (Hahnloser et al., 2003). The theory of permitted sets was further extended and developed by different authors in Curto et al. (2012, 2013). Finally, Curto and Morrison (2016) shifted attention to the study of permitted sets that can be realized as fixed point supports for a known, and uniform, external input. Like previous results, this work was largely restricted to stable fixed points of symmetric TLNs.

The nonsymmetric case, however, is considerably more interesting: asymmetric TLNs exhibit the full repertoire of nonlinear dynamic behavior, including limit cycles, quasiperiodic attractors, and chaos (Morrison, Degeratu, Itskov, & Curto, 2016). Furthermore, recent work has highlighted the importance of unstable fixed points in shaping a network’s dynamic attractors (Morrison & Curto, 2018; Morrison et al., 2019). Nevertheless, a mathematical theory connecting unstable fixed points of equation 1.1 to the structure of (W, b) has been lacking.

In this letter, we study the set of all fixed point supports, denoted $\text{FP}(W, b)$, for asymmetric, competitive W and nonnegative b . In particular, we provide two new characterizations of $\text{FP}(W, b)$: the first in terms of a sign condition (see Theorem 1), and the second in terms of domination (see Theorem 15), which for simplicity is restricted to networks with uniform b . We introduce the language of *permitted motifs* of (W, b) to refer to subsets of neurons that support a fixed point in their restricted subnetwork (all other




Figure 1: (A) A neural network with excitatory pyramidal neurons (triangles) and a background network of inhibitory interneurons (gray circles) that produce a global inhibition. The corresponding graph (right) retains only the excitatory neurons and their connections. (B) Equations for the CTLN model.

subsets are called *forbidden motifs*). We find that whether a permitted motif supports a fixed point in the full TLN depends critically on the embedding of the subnetwork inside the larger network.

In order to investigate how the fixed points are shaped by qualitative aspects of a network’s connectivity structure, we focus our attention on applications and further development of this theory in the special case of *combinatorial threshold-linear networks* (CTLNs), first introduced in Morrison et al. (2016). CTLNs are a family of competitive TLNs with uniform inputs and connectivity matrices W that are defined from simple directed graphs (see Figure 1).¹ Given a graph G and continuous parameters ε , δ and θ , the associated CTLN is the network (W, θ) where $W = W(G, \varepsilon, \delta)$ has entries:

$$W_{ij} = \begin{cases} 0 & \text{if } i = j, \\ -1 + \varepsilon & \text{if } j \rightarrow i \text{ in } G, \\ -1 - \delta & \text{if } j \not\rightarrow i \text{ in } G. \end{cases} \quad (1.2)$$

Note that $j \rightarrow i$ indicates the presence of an edge from j to i in the graph G , while $j \not\rightarrow i$ indicates the absence of such an edge. We also require that $\theta > 0$, $\delta > 0$, and $0 < \varepsilon < \frac{\delta}{\delta + 1}$. When these conditions are met, we say that the parameters are within the *legal range*.

Despite the additional constraints, CTLNs display the full range of non-linear dynamics that are observed in general TLNs (Morrison et al., 2016; Morrison & Curto, 2018), and thus provide a rich but simplified setting in which to study TLNs. Here, we focus on the theory of fixed point supports $\text{FP}(G)$ and how they can be inferred from the underlying connectivity graph G . In a companion paper (Morrison et al., 2019), we study the relationship between fixed points and dynamic attractors.

¹A graph is *simple* if it does not have loops or multiple edges between a pair of nodes.

First, we use our general results about fixed points of TLNs to develop some stronger tools that are specialized to CTLNs. These technical results then allow us to prove a series of graph rules, connecting $\text{FP}(G)$ directly to the structure of G . In particular, any conclusions about the fixed points of a CTLN that can be obtained from graph rules are automatically independent of the choice of parameters ε , δ , and θ , provided these fall within the legal range. We also study larger networks composed of smaller building block subnetworks and prove several theorems relating the fixed points and permitted motifs of the full network to those of its components. Our results demonstrate that CTLNs are surprisingly tractable and provide the foundation for a graphical calculus to infer features of a network's dynamics directly from its underlying connectivity.

The organization of this letter is as follows. In section 2, we begin by reviewing some essential background on fixed points of general TLNs, including prior results on their index and stability. Then in section 2.3, we present a new characterization of $\text{FP}(W, b)$ in terms of sign conditions (see Theorem 2). In section 3, we specialize to CTLNs and show how the sign conditions can allow us to determine fixed point supports directly from the underlying graph G . We also introduce additional tools, such as graphical domination and simply-added splits, that are later used to prove various graph rules. Sections 4 and 5, on graph rules and their extensions to composite graphs, are in some sense the heart of the letter. In particular, we identify various families of graphs and structures that yield (parameter-independent) permitted and forbidden motifs. Finally, in section 6 we introduce a more general form of domination, applicable to the broader family of TLNs, and provide a second characterization of $\text{FP}(W, b)$. This is then used to prove several of our earlier results, including Theorem 4 on graphical domination.

2 Fixed Points of TLNs

2.1 Some General Background. We refer to a network with dynamics 1.1 as a TLN (W, b) on n neurons. A *fixed point* x^* of the network is a point in the state space satisfying $\left. \frac{dx_i}{dt} \right|_{x=x^*} = 0$ for each neuron $i \in [n]$, where $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$ —in other words,

$$x_i^* = \left[\sum_{j=1}^n W_{ij} x_j^* + b_i \right]_+, \quad i = 1, \dots, n. \quad (2.1)$$

The *support* of a fixed point is the subset of active neurons,

$$\text{supp}(x^*) \stackrel{\text{def}}{=} \{i \in [n] \mid x_i^* > 0\}.$$

We typically refer to supports as subsets $\sigma \subseteq [n]$.

Definition 1. We say that a TLN (W, b) is competitive if $W_{ij} \leq 0$, $W_{ii} = 0$, and $b_i \geq 0$ for all $i, j \in [n]$.

Notice from equation 2.1 that because of the threshold nonlinearity, we must have $x_i^* \geq 0$ for each $i \in [n]$. For competitive (W, b) , it follows that whenever $x_i^* > 0$, we must have $b_i > 0$. If $b_i = 0$ for all $i \in [n]$, the activity of a competitive TLN will always decay to the fixed point $x^* = 0$. To rule this out, we also require that $b_i > 0$ for at least one neuron, ensuring that all fixed points are nontrivial (see definition 2). Furthermore, it is straightforward to check that the activity of a competitive TLN is always bounded. In particular, if $x(0) \in \prod_{i=1}^n [0, b_i]$, then $x(t) \in \prod_{i=1}^n [0, b_i]$ for all $t > 0$.

We often restrict matrices and vectors to a particular subset of neurons σ . We use the notation A_σ and b_σ to denote a matrix A and a vector b that have been truncated to include only entries with indices in σ . Furthermore, we use the notation $(A_i; b)$ to denote a matrix A whose i th column has been replaced by the vector b , as in Cramer's rule (see Lemma 1).² In the case of a restricted matrix, $((A_\sigma)_i; b_\sigma)$ denotes the matrix A_σ where the column corresponding to the index $i \in \sigma$ has been replaced by b_σ (note that this is not typically the i th column of A_σ).

Definition 2. We say that a TLN (W, b) is nondegenerate if

- $b_i > 0$ for at least one $i \in [n]$
- $\det(I - W_\sigma) \neq 0$ for each $\sigma \subseteq [n]$
- for each $\sigma \subseteq [n]$ such that $b_i > 0$ for all $i \in \sigma$, the corresponding Cramer's determinant is nonzero: $\det((I - W_\sigma)_i; b_\sigma) \neq 0$

Unless otherwise specified, we will assume all TLNs are both competitive and nondegenerate. Note that almost all networks of the form 1.1 are nondegenerate, since having a zero determinant is a highly fine-tuned condition.

If (W, b) is nondegenerate, there can be at most one fixed point per support. To see why, denote

$$x^\sigma \stackrel{\text{def}}{=} (I - W_\sigma)^{-1} b_\sigma. \quad (2.2)$$

If there exists a fixed point x^* with support σ , then $x_i^* = x_i^\sigma$ for each $i \in \sigma$ (and is zero otherwise). It follows from the definition that σ is the support of a fixed point if and only if:

1. $x_i^\sigma > 0$ for all $i \in \sigma$, and
2. $\sum_{i \in \sigma} W_{ki} x_i^\sigma + b_k \leq 0$ for all $k \notin \sigma$.

²The use of the subscript i inside $(A_i; b)$ and $((A_\sigma)_i; b_\sigma)$ has a different meaning from the subscript σ in A_σ , because it refers to replacing the i th column by b (or b_σ), as opposed to restricting the entries of A to the index set $\{\sigma\}$. This is an abuse of notation, but the meaning should always be clear from the context. We will only use the vector replacement meaning inside expressions for Cramer's determinants, such as $\det((I - W_\sigma)_i; b_\sigma)$.

(This is straightforward, but see Curto and Morrison, 2016, for more details.)

We denote the set of all fixed point supports of a TLN (W, b) as

$$\text{FP}(W, b) \stackrel{\text{def}}{=} \{\sigma \subseteq [n] \mid \sigma \text{ is the support of a fixed point}\}.$$

Finding the fixed points of a nondegenerate TLN thus reduces to finding the set of supports $\text{FP}(W, b)$. Note that because we require $b_i \geq 0$, and $b_i > 0$ for at least one i , the empty set (the support of $x^* = 0$) is never an element of $\text{FP}(W, b)$.

Since condition 1 above depends on only (W_σ, b_σ) , a necessary condition for $\sigma \in \text{FP}(W, b)$ is that $\sigma \in \text{FP}(W_\sigma, b_\sigma)$. Such a fixed point survives the addition of other nodes $k \notin \sigma$ precisely when condition 2 is satisfied. Note that only the survival of the fixed point depends on the rest of the network; the actual values of the x_i^σ (for $i \in \sigma$) cannot change. For this reason, it makes sense to distinguish the subsets σ that support a fixed point on the restricted networks (W_σ, b_σ) , irrespective of whether these fixed points survive to the full network. In particular, if $\sigma \notin \text{FP}(W_\sigma, b_\sigma)$, we are guaranteed that $\sigma \notin \text{FP}(W, b)$.

Definition 3 (*permitted and forbidden motifs*). Let (W, b) be a TLN on n neurons. We say that $\sigma \subseteq [n]$ is a permitted motif of the network if $\sigma \in \text{FP}(W_\sigma, b_\sigma)$. Otherwise, we say that σ is a forbidden motif.

2.2 Index and Parity. For each TLN fixed point, labeled by its support $\sigma \in \text{FP}(W, b)$, we define the index as

$$\text{idx}(\sigma) \stackrel{\text{def}}{=} \text{sgn det}(I - W_\sigma).$$

Since we assume our TLNs are nondegenerate, $\text{det}(I - W_\sigma) \neq 0$ and thus $\text{idx}(\sigma) \in \{\pm 1\}$. Moreover, if σ is the support of a stable fixed point, then the eigenvalues of $-I + W_\sigma$ must all have negative real part, and so those of $I - W_\sigma$ all have positive real part. This implies that $\text{idx}(\sigma) = +1$ for all stable fixed points.

The following theorem, given in Morrison et al. (2019), indicates that fixed points with index $+1$ and -1 are almost perfectly balanced. It also tells us that the parity of the total number of fixed points is always odd.

Theorem 1 (*parity; Morrison et al., 2019*). Let (W, b) be a TLN. Then

$$\sum_{\sigma \in \text{FP}(W, b)} \text{idx}(\sigma) = +1.$$

In particular, the total number of fixed points $|\text{FP}(W, b)|$ is always odd.

As an immediate corollary, we obtain an upper bound on the number of stable fixed points, which all have index +1. Here we also use the fact that $|\text{FP}(W, b)| \leq 2^n - 1$, which is the number of nonempty subsets of $[n]$.

Corollary 1. *The number of stable fixed points in a TLN on n neurons is at most 2^{n-1} .*

2.3 Sign Conditions. In this section, we provide a new characterization of fixed point supports of TLNs via the signs of particular Cramer's determinants. Recall Cramer's rule:

Lemma 1. *Let A be an $n \times n$ matrix with $\det A \neq 0$, and consider the linear system $Ax = b$. This has a unique solution, $x = A^{-1}b$. The entries of x can be expressed as*

$$x_i = \frac{\det(A_i; b)}{\det A},$$

where $(A_i; b)$ is the matrix obtained from A by replacing the i th column with b .

Indeed, it follows directly from Cramer's rule that a fixed point of a TLN (W, b) with support σ has values

$$x_i^\sigma = \frac{\det((I - W_\sigma)_i; b_\sigma)}{\det(I - W_\sigma)}, \quad \text{for } i \in \sigma. \quad (2.3)$$

For any $\sigma \subseteq [n]$, we are thus motivated to define

$$s_i^\sigma \stackrel{\text{def}}{=} \det((I - W_{\sigma \cup \{i\}})_i; b_{\sigma \cup \{i\}}), \quad \text{for each } i \in [n]. \quad (2.4)$$

Note that because we consider only TLNs that are nondegenerate (see definition 2), we can assume all of Cramer's determinants are nonzero, and thus $s_i^\sigma \neq 0$ for all $i \in [n]$ and $\sigma \subseteq [n]$. Moreover, it follows directly from the definition that for any $i \in [n]$,

$$s_i^\sigma = s_i^{\sigma \cup \{i\}}. \quad (2.5)$$

Note that for the empty set, we have $s_i^\emptyset = s_i^{\{i\}} = b_i$.

It turns out that fixed point supports of (competitive, nondegenerate) TLNs can be fully determined from the signs of the s_i^σ , yielding our first characterization of $\text{FP}(W, b)$. Recall that σ is a permitted motif of (W, b) if $\sigma \in \text{FP}(W_\sigma, b_\sigma)$.

Theorem 2 (sign conditions). *Let (W, b) be a TLN on n neurons. For any nonempty $\sigma \subseteq [n]$,*

$$\sigma \text{ is a permitted motif} \iff \text{sgn } s_i^\sigma = \text{sgn } s_j^\sigma \text{ for all } i, j \in \sigma.$$

When σ is permitted, $\text{sgn } s_i^\sigma = \text{sgn } \det(I - W_\sigma) = \text{idx}(\sigma)$ for all $i \in \sigma$. Furthermore,

$$\sigma \in \text{FP}(W, b) \Leftrightarrow \text{sgn } s_j^\sigma = \text{sgn } s_i^\sigma = -\text{sgn } s_k^\sigma \text{ for all } i, j \in \sigma, k \notin \sigma.$$

Before we prove Theorem 2, we need the following lemma, which gives a useful identity for computing s_k^σ values:

Lemma 2. *Let (W, b) be a TLN on n neurons, and let $\sigma \subseteq [n]$. Then*

$$s_k^\sigma = \sum_{i \in \sigma} W_{ki} s_i^\sigma + b_k \det(I - W_\sigma) \text{ for any } k \in [n]. \quad (2.6)$$

Proof. For $k \in \sigma$, it follows from the definition of x^σ , equation 2.2, that $x^\sigma = W_\sigma x^\sigma + b_\sigma$, and thus

$$x_k^\sigma = \sum_{i \in \sigma} W_{ki} x_i^\sigma + b_k.$$

Using equation 2.3, this yields $s_k^\sigma = \sum_{i \in \sigma} W_{ki} s_i^\sigma + b_k \det(I - W_\sigma)$, as desired. Next, we consider $k \notin \sigma$ and compute

$$s_k^\sigma = \det((I - W_{\sigma \cup \{k\}})_k; b_{\sigma \cup \{k\}}) = \det\left(\frac{I - W_\sigma}{-W_{k1} \cdots -W_{k,k-1}} \middle| \frac{b_\sigma}{b_k}\right).$$

Applying the Laplace expansion for the determinant along the k th row, we obtain

$$\begin{aligned} s_k^\sigma &= \sum_{i \in \sigma} (-1)^{i+k} (-W_{ki}) (-1)^{(k-1)-i} \det((I - W_\sigma)_i; b_\sigma) + b_k \det(I - W_\sigma) \\ &= \sum_{i \in \sigma} W_{ki} s_i^\sigma + b_k \det(I - W_\sigma), \end{aligned}$$

which completes the proof. \square

We are now ready to prove Theorem 2:

Proof of Theorem 2 (sign conditions). Recall that $\sigma \in \text{FP}(W_\sigma, b_\sigma)$ if and only if $x_i^\sigma > 0$ for each $i \in \sigma$ (see equation 2.2 and subsequent remarks). By Cramer's rule, we have $x_i^\sigma = \frac{s_i^\sigma}{\det(I - W_\sigma)}$ (see equation 2.3). Now suppose $\sigma \in \text{FP}(W_\sigma, b_\sigma)$. Since $x_i^\sigma > 0$ for each $i \in \sigma$, we must have $\text{sgn } s_i^\sigma = \text{sgn } s_j^\sigma = \text{sgn } \det(I - W_\sigma) = \text{idx}(\sigma)$ for all $i, j \in \sigma$.

For the reverse implication, suppose $\text{sgn } s_i^\sigma = \text{sgn } s_j^\sigma$ for all $i, j \in \sigma$. This immediately implies that all the x_i^σ for $i \in \sigma$ have the same sign, but we must show this sign is positive. First, we show that $b_i > 0$ for all $i \in \sigma$ (in

competitive TLNs, we always have $b_i \geq 0$). To see this, suppose there exists a $j \in \sigma$ such that $b_j = 0$. Then by equation 2.6 we would have

$$s_j^\sigma = \sum_{i \in \sigma} W_{ji} s_i^\sigma + b_j \det(I - W_\sigma) = \sum_{i \in \sigma \setminus j} W_{ji} s_i^\sigma,$$

using the fact that $W_{jj} = 0$. Since $W_{ji} < 0$ for all $i \neq j$, the above equality contradicts the assumption that all the signs of the s_i^σ are equal for $i \in \sigma$. We can thus conclude that $b_i > 0$ for all $i \in \sigma$. This in turn ensures that $x_i^\sigma > 0$ for at least one (and thus all) $i \in \sigma$, because by definition $(I - W_\sigma)x_i^\sigma = b_\sigma$, and all entries of $(I - W_\sigma)$ are nonnegative. Hence, $\sigma \in \text{FP}(W_\sigma, b_\sigma)$, completing the proof of the first part of the theorem.

To prove the second part, recall that $\sigma \in \text{FP}(W, b)$ precisely when $\sigma \in \text{FP}(W_\sigma, b_\sigma)$ and σ satisfies fixed point condition 2: $\sum_{i \in \sigma} W_{ki} x_i^\sigma + b_k \leq 0$ for all $k \notin \sigma$. Using equation (2.6) again and dividing by $\det(I - W_\sigma)$, we find that

$$\frac{s_k^\sigma}{\det(I - W_\sigma)} = \sum_{i \in \sigma} W_{ki} x_i^\sigma + b_k. \quad (2.7)$$

Since the network is nondegenerate, $\det(I - W_\sigma) \neq 0$ and $s_k^\sigma \neq 0$, and so condition 2 is equivalent to the sign condition: $\text{sgn } s_k^\sigma = -\text{sgn } \det(I - W_\sigma)$. Putting this together with the above sign conditions for $\sigma \in \text{FP}(W_\sigma, b_\sigma)$, we see that $\sigma \in \text{FP}(W, b)$ if and only if $\text{sgn } s_i^\sigma = \text{sgn } s_j^\sigma = -\text{sgn } s_k^\sigma$ for all $i, j \in \sigma$, $k \notin \sigma$. \square

In the following example, we show how to use Theorem 2 to find $\text{FP}(W, b)$ for a TLN of size 2 with a uniform external input b .

Example 1. Consider a TLN with $W = \begin{pmatrix} 0 & W_{12} \\ W_{21} & 0 \end{pmatrix}$, for some $W_{12}, W_{21} < 0$, and external input $b = \mathbf{1}$. Recall that the empty set is never a fixed point support of a competitive TLN with a positive input, so we restrict to considering nonempty subsets of $\{1, 2\}$. For $\sigma = \{1\}$, we see that $s_1^\sigma = b_1 = 1$, while

$$s_2^\sigma = \det((I - W_{\sigma \cup 2})_2; b) = \det \begin{pmatrix} 1 & 1 \\ -W_{21} & 1 \end{pmatrix} = 1 + W_{21}.$$

In particular, $\text{sgn } s_2^\sigma = -\text{sgn } s_1^\sigma$ precisely when $W_{21} < -1$. Thus, by Theorem 2,

$$\{1\} \in \text{FP}(W, b) \Leftrightarrow W_{21} < -1.$$

By a similar argument, we have

$$\{2\} \in \text{FP}(W, b) \Leftrightarrow W_{12} < -1.$$

Finally, consider $\sigma = \{1, 2\}$. We obtain $s_1^\sigma = 1 + W_{12}$ and $s_2^\sigma = 1 + W_{21}$. By Theorem 2, $\sigma \in \text{FP}(W, b)$ precisely when $\text{sgn}(1 + W_{12}) = \text{sgn}(1 + W_{21})$, and so

$$\{1, 2\} \in \text{FP}(W, b) \Leftrightarrow \begin{cases} W_{12} > -1 \text{ and } W_{21} > -1, \text{ or} \\ W_{12} < -1 \text{ and } W_{21} < -1 \end{cases}.$$

Note that in the case where W is a CTLN, the condition $W_{21} < -1$ corresponds to $1 \not\rightarrow 2$ in the associated graph G , while $W_{12} < -1$ indicates $2 \not\rightarrow 1$ in G . So the last condition reduces to $\{1, 2\} \in \text{FP}(G|_\sigma)$ if and only if nodes 1 and 2 are either bidirectionally connected or disconnected in G .

Since the values of s_i^σ for $i \in \sigma$ depend only on (W_σ, b_σ) , while the values s_k^σ for $k \notin \sigma$ depend only on $(W_{\sigma \cup \{k\}}, b_{\sigma \cup \{k\}})$, we immediately have the following useful corollary:

Corollary 2. *Let (W, b) be a TLN on n neurons, and let $\sigma \subseteq [n]$. The following are equivalent:*

1. $\sigma \in \text{FP}(W, b)$
2. $\sigma \in \text{FP}(W_\tau, b_\tau)$ for all $\sigma \subseteq \tau \subseteq [n]$.
3. $\sigma \in \text{FP}(W_\sigma, b_\sigma)$ and $\sigma \in \text{FP}(W_{\sigma \cup \{k\}}, b_{\sigma \cup \{k\}})$ for all $k \notin \sigma$
4. $\sigma \in \text{FP}(W_{\sigma \cup \{k\}}, b_{\sigma \cup \{k\}})$ for all $k \notin \sigma$

Theorem 2 also gives a relationship between the indices of fixed points whose supports differ by only one neuron.

Lemma 3 (alternation). *Let (W, b) be a TLN. If $\sigma, \sigma \cup \{k\} \in \text{FP}(W, b)$, for $k \notin \sigma$, are both fixed point supports, then*

$$\text{idx}(\sigma \cup \{k\}) = -\text{idx}(\sigma).$$

Proof. If $\sigma, \sigma \cup \{k\} \in \text{FP}(W, b)$, then by Theorem 2, we have $\text{idx}(\sigma) = \text{sgn } s_i^\sigma = -\text{sgn } s_k^\sigma$ for any $i \in \sigma$. Recalling that $s_k^{\sigma \cup \{k\}} = s_k^\sigma$, we see that $\text{idx}(\sigma \cup \{k\}) = -\text{idx}(\sigma)$, as desired. \square

Corollary 3. *If $\sigma \in \text{FP}(W, b)$ is the support of a stable fixed point, then there is no other stable fixed point with support $\sigma, \sigma \setminus k$, or $\sigma \cup \{k\}$ for any $k \in [n]$.*

3 Fixed Points of CTLNs

When (W, b) comes from a CTLN with graph G and parameters ε, δ , and θ , so that $W = W(G, \varepsilon, \delta)$ and $b = \theta \mathbf{1}$, we use the notation

$$\text{FP}(G) = \text{FP}(G, \varepsilon, \delta) \stackrel{\text{def}}{=} \text{FP}(W(G, \varepsilon, \delta), \theta \mathbf{1}).$$

Note that we always suppress θ from the notation, since it can easily be seen that the value of θ does not affect the set of fixed point supports, so long as $\theta > 0$. (Different θ values merely rescale the fixed point values.) In addition, we typically suppress the ε and δ dependence as well, using the simpler notation $\text{FP}(G)$ to denote the set of fixed point supports when ε and δ are understood to be fixed.

For some graphs, $\text{FP}(G) = \text{FP}(G, \varepsilon, \delta)$ does indeed depend on the choice of ε and δ , but our theoretical results almost always deal with aspects of $\text{FP}(G)$ that are independent of the choice of parameters, provided these lie within the legal range.³ For example, a graph rule could tell us about certain fixed point supports that can be ruled out and are thus not contained in $\text{FP}(G)$. Such a conclusion is parameter independent, even if there are other supports in $\text{FP}(G)$ whose presence depends on parameters.

Recall that σ is a permitted motif if $\sigma \in \text{FP}(W_\sigma, b_\sigma)$ and is forbidden otherwise. We use the same language for CTLNs: σ is a *permitted motif* if $\sigma \in \text{FP}(G|_\sigma)$, where $G|_\sigma$ is the induced subgraph obtained by restricting vertices and edges to the vertex set σ . If $\sigma \notin \text{FP}(G|_\sigma)$, we say that σ is a *forbidden motif*. Note that this may depend on the choice of parameters ε, δ .

The results in this section provide a technical foundation for the graph rules and building block rules that we present in sections 4 and 5, respectively. Before moving on to the main content, we pause briefly to provide some simple bounds on the total activity at fixed points of CTLNs, irrespective of the support. These bounds, like the actual fixed point values, do depend on ε, δ , and θ , even when the supports $\text{FP}(G)$ do not.

Lemma 4. *If x^* is a fixed point of a CTLN on n nodes, with parameters ε, δ , and θ , then*

$$\frac{\theta}{1 + \delta} < \sum_{i=1}^n x_i^* < \frac{\theta}{1 - \varepsilon}.$$

Proof. Let σ be the support of x^* . Then $(I - W_\sigma)x_\sigma^* = \theta \mathbf{1}_\sigma$, and so $\theta = R_j \cdot x_\sigma^*$, where R_j is the j th row vector of $I - W_\sigma$. Since all entries of R_j and x_σ^* are positive and the off-diagonal entries of $I - W_\sigma$ are all at least $1 - \varepsilon$ and at most $1 + \delta$, it follows that

$$(1 - \varepsilon) \sum_{i \in \sigma} x_i^* + \varepsilon x_j^* \leq \theta \leq (1 + \delta) \sum_{i \in \sigma} x_i^* - \delta x_j^*$$

for any $j \in \sigma$. Since $x_j^* > 0$, we have $(1 - \varepsilon) \sum_{i \in \sigma} x_i^* < \theta < (1 + \delta) \sum_{i \in \sigma} x_i^*$, which implies the desired result. \square

³For any result that does depend on ε and δ , we indicate this explicitly.




Figure 2: Graph with s_i^σ values for all nonempty σ . (A) Graph with $\text{FP}(G, \varepsilon, \delta)$. (B1–B7) In each panel, the nodes in the gray-shaded region comprise σ , and each node is labeled with its s_i^σ value.

3.1 Sign Rules for CTLNs. Recall from section 2.3 that Theorem 2 gives sign conditions for when a subset σ is a permitted motif or a fixed point support in a general TLN (W, b) . We illustrate this result for the special case of CTLNs in the following example.

Example 2. Let G be the graph in Figure 2A. We will use the sign conditions to find the permitted motifs and $\text{FP}(G) = \text{FP}(G, \varepsilon, \delta)$. Recall from section 2.1 that the empty set is never the support of a fixed point, and so we restrict to analyzing the nonempty subsets of $\{1, 2, 3\}$. Each panel in Figure 2 shows a choice of σ (gray-shaded regions) and the values of $s_i^\sigma = \det((I - W_{\sigma \cup i})_i; \theta \mathbf{1})$ for all $i \in \{1, 2, 3\}$. For example, in panel B1, $s_1^{(1)} = \det(\theta) = \theta$ and $s_2^{(1)} = \det \begin{pmatrix} 1 & \theta \\ 1 - \varepsilon & \theta \end{pmatrix} = \varepsilon\theta$.

Observe that every singleton $\{i\}$ is a permitted motif, as are $\{1, 2\}$, $\{2, 3\}$, and $\{1, 2, 3\}$ since these subsets satisfy $\text{sgn } s_i^\sigma = \text{sgn } s_j^\sigma$ for all $i, j \in \sigma$. A permitted motif survives as a fixed point of G precisely when the external nodes all have the opposite sign for s_k^σ . Thus, the only singleton fixed point support is $\{3\}$ since it satisfies $\text{sgn } s_i^{(3)} = -\text{sgn } s_3^{(3)}$ for $i = 1, 2$ (see panel B3 in Figure 2). Continuing this analysis, we find that $\text{FP}(G) = \{3, 12, 123\}$. Furthermore, since the signs of s_i^σ are constant across $\varepsilon, \delta > 0$ for each $i \in \{1, 2, 3\}$ and $\sigma \subseteq \{1, 2, 3\}$, we see that $\text{FP}(G)$ is in fact parameter independent.

Figure 18 in appendix section A.1 shows all directed graphs of size $n \leq 3$ labeled with the full support $s_i^{[n]}$ values for each node $i \in [n]$. From this, we see that 6 of the 16 directed graphs of size 3 are permitted motifs, and the remainder are forbidden, independent of the choice of parameters.

3.2 Simply-Added Splits. In this section, we introduce the concept of *simply-added splits*, whereby the vertices of a graph are partitioned in a




Figure 3: In this graph, ω is simply-added to τ , and thus each $k \in \omega$ either sends all possible edges to τ or no edges. There is no constraint on the edges within τ , within ω , or from τ to ω .

special way that allows us to easily compute certain s_i^σ values. This kind of structure will play a prominent role in section 5, where we build larger graphs from component subgraphs. In particular, simply-added splits are essential to our ability to relate $\text{FP}(G)$ for such a composite graph to the $\text{FP}(G_i)$ of its components. They also provide an additional tool for proving some of the graph rules in section 4.

Definition 4 (*simply-added*). Let G be a graph on n nodes, $\tau \subset [n]$ nonempty, and $k \notin \tau$. We say that k is a projector onto τ if $k \rightarrow i$ for all $i \in \tau$. We say that k is a nonprojector onto τ if $k \not\rightarrow i$ for all $i \in \tau$. For any nonempty $\tau, \omega \subset [n]$ such that $\tau \cap \omega = \emptyset$, we say ω is simply-added to τ if for each $k \in \omega$, k is either a projector or a nonprojector onto τ (see Figure 3). In this case, we say that the (τ, ω) is a simply-added split of the subgraph $G|_\sigma$, for $\sigma = \tau \cup \omega$.

The following theorem shows that the s_i^σ values factor nicely when a set of neurons is simply-added to a subgraph.

Theorem 3 (*simply-added*). Let G be a graph on n nodes, and let $\tau, \omega \subset [n]$ be such that ω is simply-added to τ . For $\sigma = \tau \cup \omega$,

$$s_i^\sigma = \frac{1}{\theta} s_i^\omega s_i^\tau = \alpha s_i^\tau \quad \text{for each } i \in \tau,$$

where $\alpha = \frac{1}{\theta} s_i^\omega$ has the same value for every $i \in \tau$.

More generally, for any $\sigma \subseteq \tau \cup \omega$,

$$s_i^\sigma = \frac{1}{\theta} s_i^{\sigma \cap \omega} s_i^{\sigma \cap \tau} = \alpha s_i^{\sigma \cap \tau} \quad \text{for each } i \in \tau,$$

where $\alpha = \frac{1}{\theta} s_i^{\sigma \cap \omega}$.

Recall, using equation 2.5 that $s_i^\emptyset = \theta$, since for CTLNs $b_i = \theta$ for each $i \in [n]$. With this convention, Theorem 3 holds even if $\sigma \cap \tau$ or $\sigma \cap \omega$ is empty.

To prove Theorem 3, we use the standard formula for the determinant of a 2×2 block matrix,

$$\det \begin{bmatrix} A & B \\ C & D \end{bmatrix} = \det(A)\det(D - CA^{-1}B), \quad (3.1)$$

which applies as long as A is invertible.

Proof of Theorem 3. Since ω is simply-added to τ , each $k \in \omega$ is either a projector or nonprojector onto τ . Thus, for each $k \in \omega$, $W_{ik} = -1 + c_k$ for all $i \in \tau$, where $c_k = \varepsilon$ if k is a projector or $c_k = -\delta$ if k is nonprojector.

First consider $\sigma = \tau \cup \omega$. For any $i \in \tau$, we compute:

$$\begin{aligned} s_i^\sigma &= \theta \det((I - W_\sigma)_i; \mathbf{1}) = \theta \det \left(\begin{array}{c|ccc} 1 & 1 - c_{k_1} & \cdots & 1 - c_{k_\ell} \\ I - W_{\tau \setminus \{i\}} & \vdots & & \vdots \\ 1 & 1 - c_{k_1} & \cdots & 1 - c_{k_\ell} \\ \hline * & 1 & 1 - c_{k_1} & \cdots & 1 - c_{k_\ell} \\ * & \mathbf{1} & I - W_\omega & & \end{array} \right) \\ &= \theta \det \left(\begin{array}{c|ccc} 1 & 0 & \cdots & 0 \\ I - W_{\tau \setminus \{i\}} & \vdots & & \vdots \\ 1 & 0 & \cdots & 0 \\ \hline * & 1 & 0 & \cdots & 0 \\ * & \mathbf{1} & I - W_\omega - A & & \end{array} \right) \\ &= \theta \det((I - W_\tau)_i; \mathbf{1}) \det(I - W_\omega - A), \end{aligned}$$

where A consists of the modifications to the $I - W_\omega$ block that arise from clearing out the top right block using the column of 1s.

Now observe that $\theta \det((I - W_\tau)_i; \mathbf{1}) = s_i^\tau$, while

$$\begin{aligned} \det(I - W_\omega - A) &= \det \left(\begin{array}{c|ccc} 1 & 0 & \cdots & 0 \\ \mathbf{1} & I - W_\omega - A & & \end{array} \right) \\ &= \det \left(\begin{array}{c|ccc} 1 & 1 - c_{k_1} & \cdots & 1 - c_{k_\ell} \\ \mathbf{1} & I - W_\omega & & \end{array} \right) \\ &= \det((I - W_{\omega \cup \{i\}})_i; \mathbf{1}) = \frac{1}{\theta} s_i^{\omega \cup \{i\}} = \frac{1}{\theta} s_i^\omega. \end{aligned}$$

We therefore have $s_i^\sigma = \frac{1}{\theta} s_i^\tau s_i^\omega$ and see that s_i^σ has the same value for each $i \in \tau$.

Now consider $\sigma \subseteq \tau \cup \omega$. Note that $\sigma = (\sigma \cap \tau) \cup (\sigma \cap \omega)$, where $\sigma \cap \omega$ is simply-added to $\sigma \cap \tau$ (since ω is simply-added to τ). To compute s_i^σ for $i \in \tau$, we consider two cases. First, suppose $i \in \sigma \cap \tau$. In this case, applying

the previous formula to $\sigma = (\sigma \cap \tau) \cup (\sigma \cap \omega)$ immediately yields the desired result. Now consider $i \in \tau \setminus \sigma$. In this case, we can apply the previous formula to $\sigma \cup i = ((\sigma \cap \tau) \cup i) \cup (\sigma \cap \omega)$, and so

$$s_i^\sigma = s_i^{\sigma \cup i} = \frac{1}{\theta} s_i^{\sigma \cap \omega} s_i^{(\sigma \cap \tau) \cup i} = \frac{1}{\theta} s_i^{\sigma \cap \omega} s_i^{\sigma \cap \tau},$$

where the first and last equalities use equation 2.5. Thus, the desired formula for s_i^σ holds for all $i \in \tau$. \square

Note that if $\omega = \{k\}$, where k is a projector onto τ , then $s_i^\omega = \varepsilon\theta$ for each $i \in \tau$. On the other hand, if k is a nonprojector onto τ , then $s_i^\omega = -\delta\theta$ for each $i \in \tau$. To see this, observe that $s_i^{\{k\}} = s_i^{\{i,k\}} = (1 + W_{ik})\theta$, which evaluates to $\varepsilon\theta$ if $k \rightarrow i$, and $-\delta\theta$ for $k \not\rightarrow i$. This gives the following useful corollary.

Corollary 4. *Suppose $\sigma = \tau \cup \{k\}$, where k is simply-added onto τ . If k is a projector, then $s_i^\sigma = \varepsilon s_i^\tau$ for each $i \in \tau$. If k is a nonprojector, then $s_i^\sigma = -\delta s_i^\tau$ for each $i \in \tau$.*

Example 3. Let G be the graph in Figure 2A. For $\sigma = \{1, 2, 3\}$, we see that σ has a simply-added split $\sigma = \tau \cup \omega$ where $\tau = \{1, 2\}$ and $\omega = \{3\}$ is a nonprojector onto τ . Applying Corollary 4, we have $s_i^\sigma = -\delta s_i^\tau$ for $i = 1, 2$. We see from panel B4 in Figure 2 that $s_i^\tau = \varepsilon\theta$, and panel B7 shows $s_i^\sigma = -\delta\varepsilon\theta$ for $i = 1, 2$. Furthermore, σ actually has a second simply-added split: $\sigma = \tau \cup \omega$, where $\tau = \{2, 3\}$ and $\omega = \{1\}$ is a projector onto τ . In this case, Corollary 4 guarantees that $s_i^\sigma = \varepsilon s_i^\tau$ for $i = 2, 3$. Since $s_i^\tau = -\delta\theta$ for $i = 2, 3$ (see panel B6 of Figure 2), we see that this second simply-added split gives a consistent value for s_2^σ .

3.3 Graphical Domination. One of the most important tools we will use to prove the graph rules in section 4 is the concept of graphical domination. This refers to certain “domination” relationships between vertices in the graph of a CTLN that can be identified by examining the graph alone (e.g., without computing s_i^σ values). The presence of such a relationship within a subgraph $G|_\sigma$ is sufficient to guarantee that σ is a forbidden motif. Graphical domination can also be used to determine whether a permitted motif survives as a fixed point support in a larger network. These facts are collected in Theorem 4. Because they rely only on the graph structure, any results about a CTLN obtained from graphical domination are automatically parameter independent, within the legal range.

Although we will make frequent use of Theorem 4 in sections 4 and 5, we postpone the proof until section 6.3. This is because graphical domination is a special case of a more general notion of domination, which we introduce in section 6.

Definition 5. *We say that k graphically dominates j with respect to σ if $\sigma \cap \{j, k\} \neq \emptyset$ and the following three conditions all hold:*




Figure 4: The three cases of graphical domination in Theorem 4. In each panel, k graphically dominates j with respect to σ (the outermost shaded region). The inner shaded regions illustrate the subsets of nodes that send edges to j and k . Note that the vertices sending edges to j are a subset of those sending edges to k , but this containment need not be strict. Dashed arrows indicate optional edges between j and k .

1. For each $i \in \sigma \setminus \{j, k\}$, if $i \rightarrow j$ then $i \rightarrow k$.
2. If $j \in \sigma$, then $j \rightarrow k$.
3. If $k \in \sigma$, then $k \not\rightarrow j$.

Theorem 4 (graphical domination). Suppose k graphically dominates j with respect to σ (see Figure 4). Then the following statements all hold:

1. If $j, k \in \sigma$, then $\sigma \notin FP(G|_\sigma)$, and so $\sigma \notin FP(G)$. (inside-in domination)
2. If $j \in \sigma$ and $k \notin \sigma$, then $\sigma \notin FP(G|_{\sigma \cup \{k\}})$, and so $\sigma \notin FP(G)$. (outside-in domination)
3. If $j \notin \sigma, k \in \sigma$, and $\sigma \in FP(G|_\sigma)$, then $\sigma \in FP(G|_{\sigma \cup \{j\}})$. (inside-out domination)

Graphical domination is a special case of a more general notion of domination, which applies to any TLN. We discuss general domination in section 6, and show in Theorem 15 that it gives us a complete characterization of fixed point supports, similar to the sign conditions in Theorem 2. In section 6.3, we use general domination to prove Theorem 4.

Part 1 of Theorem 4 tells us that if there is any graphical domination inside σ , then σ is not a permitted motif and therefore cannot be a fixed point support in $FP(G)$. Part 2 tells us that if there is any node $j \in \sigma$ that is dominated by a k outside σ , then σ does not survive as a fixed point support of G , irrespective of whether σ is a permitted motif. Finally, part 3 implies that if for each $j \notin \sigma$ there exists a k inside σ that dominates j , then σ is guaranteed to survive as a fixed point support in $FP(G)$, provided σ is a permitted motif.

It is important to note that any results obtained via graphical domination hold for all values of ε, δ within the legal range. Thus, if σ is forbidden by graphical domination, it is forbidden independent of parameters; similarly, if σ survives (or does not survive) due to graphical domination, this fact is

also parameter independent. This is not necessarily the case for forbidden motifs that do not have graphical domination (see appendix section A.2 for examples).

We will see in section 5 that when a motif is forbidden by graphical domination, this feature gets inherited to larger graphs that have this motif as their skeleton. This follows from the fact that graphical domination interplays nicely with simply-added splits.⁴ For this reason, it is useful to define a stronger notion of forbidden motif:

Definition 6 (*strongly forbidden*). Consider a CTLN on n nodes. We say that $\sigma \subseteq [n]$ is strongly forbidden if there exist $j, k \in \sigma$ such that k graphically dominates j with respect to σ .

3.4 Uniform In-Degree. Here we introduce uniform in-degree graphs, which comprise a large and interesting family of permitted motifs. This family is particularly nice because the survival rules are parameter independent and can be easily checked directly from the graph (see Theorem 5).

Definition 7. We say that σ has uniform in-degree d if every $i \in \sigma$ has in-degree $d_i^{\text{in}} = d$ within $G|_\sigma$.

Note that the subgraph $G|_\sigma$ could have uniform in-degree, but the nodes of σ may have different degrees with respect to the full graph G .

There exist graphs with uniform in-degree d for any $0 \leq d \leq |\sigma| - 1$. When $d = 0$, the graph is necessarily an *independent set*, that is, a collection of nodes with no edges between them; when $d = |\sigma| - 1$, it is a *clique*, that is, every pair of nodes is connected by a bidirectional edge. For in-between values of d , however, there are several distinct possibilities. For example, if $|\sigma| = 4$ and $d = 1$, the subgraph $G|_\sigma$ could be a 4-cycle, a pair of 2-cliques, or a 3-cycle with a single outgoing edge to a fourth node (we call this last graph the “tadpole”). Notice that the tadpole graph illustrates that a uniform in-degree subgraph need not be cyclically symmetric and need not have uniform out-degree. Figure 5 illustrates several examples of uniform in-degree graphs.

Theorem 5 (*uniform in-degree*). Let σ have uniform in-degree d . For $k \notin \sigma$, let $d_k \stackrel{\text{def}}{=} |\{i \in \sigma \mid i \rightarrow k\}|$ be the number of edges k receives from σ . Then

$$\sigma \in \text{FP}(G|_{\sigma \cup \{k\}}) \Leftrightarrow d_k \leq d.$$

Furthermore, if $|\sigma| > 1$ and $d < |\sigma|/2$, then the fixed point is unstable. If $d = |\sigma| - 1$ (i.e., if σ is a clique), then the fixed point is stable.

⁴Suppose ω is simply-added to τ . If $j, k \in \tau$, and k graphically dominates j with respect to τ , then k also graphically dominates j with respect to $\tau \cup \omega$. This is because all edges from ω to j have corresponding edges from ω to k .




Figure 5: Examples of uniform in-degree graphs. (A1–A4) All graphs with uniform in-degree on $n = 3$ vertices. (B1–B4) Some examples with $n = 4$ and $n = 5$ vertices.

The proof of Theorem 5 requires general domination and is thus postponed to section 6.4. Note that this result, like graphical domination, is parameter independent.

Next we show how the eigenvalues of a uniform in-degree subnetwork $I - W_\tau$ are inherited when a subset ω is simply-added to τ . Recall that in the proof of Theorem 3 (simply-added), the determinant calculation was simplified by the insertion of the all-ones vector, which occurred in the expressions $s_i^\sigma = \theta \det((I - W_\sigma)_i; \mathbf{1})$. When τ has uniform in-degree, we can apply a similar trick to compute the eigenvalues of $I - W_\sigma$, where $\sigma = \tau \cup \omega$.

Lemma 5. *Let σ have a simply-added split $\tau \cup \omega$, where τ is uniform in-degree. Let R_τ be the row sum of $I - W_\tau$. Note that this is the top (Perron-Frobenius) eigenvalue of $I - W_\tau$. Then*

$$\text{eig}(I - W_\sigma) \supset \text{eig}(I - W_\tau) \setminus R_\tau.$$

So all the eigenvalues of τ get inherited, except possibly the top one R_τ . In particular, if τ is unstable, then σ is also unstable.

Proof. To calculate $\text{eig}(I - W_\sigma)$, observe that

$$\det(I - W_\sigma - \lambda I) = \det \left(\begin{array}{ccc|ccc} & & & \alpha_1 & \dots & \alpha_{|\omega|} \\ & & & \vdots & & \vdots \\ I - W_\tau - \lambda I & & & \alpha_1 & \dots & \alpha_{|\omega|} \\ \hline * & & & I - W_\omega - \lambda I & & \end{array} \right),$$

where $\alpha_i = 1 - \varepsilon$ if i is a projector onto τ , and $\alpha_i = 1 + \delta$ otherwise. Let

$$A = ((I - W_\tau - \lambda I)_{|\tau|}; (R_\tau - \lambda)\mathbf{1})$$

be the matrix obtained from $I - W_\tau - \lambda I$ by replacing the last column with a column of $R_\tau - \lambda$. Note that $\det(I - W_\tau - \lambda I) = \det(A)$ since A can be obtained from $I - W_\tau - \lambda I$ by adding all the columns of the matrix to its last column. Performing these same column operations on the full matrix yields

$$\det(I - W_\sigma - \lambda I) = \det \left(\begin{array}{c|ccc} & \alpha_1 & \dots & \alpha_{|\omega|} \\ A & \vdots & & \vdots \\ & \alpha_1 & \dots & \alpha_{|\omega|} \\ \hline * & & I - W_\omega - \lambda I & \end{array} \right).$$

Using the last column of A (scaled by $\frac{\alpha_i}{R_\tau - \lambda}$), we can clear out the top right of the matrix to obtain

$$\det(I - W_\sigma - \lambda I) = \det \left(\begin{array}{c|c} A & 0 \\ \hline * & B \end{array} \right),$$

where $B = \frac{1}{R_\tau - \lambda} C$ for some matrix C whose entries are all polynomial in λ . Note that $\det B = \frac{1}{(R_\tau - \lambda)^{|\omega|}} \det C$, and so

$$\det(I - W_\sigma - \lambda I) = \det(A) \det(B) = \det(I - W_\tau - \lambda I) \frac{\det(C)}{(R_\tau - \lambda)^{|\omega|}}.$$

Thus, all the roots of $\det(I - W_\tau - \lambda I)$ must also be roots of $\det(I - W_\sigma - \lambda I)$ except possibly $\lambda = R_\tau$, and so $\text{eig}(I - W_\sigma) \supset \text{eig}(I - W_\tau) \setminus R_\tau$ as desired. \square

This result will be useful in section 5.1 in the context of composite graphs (see Proposition 1).

4 Graph Rules for CTLNs

In this section, we prove a variety of graph rules characterizing fixed point supports of CTLNs in terms of the underlying graph G . These are truly “graph rules” in that they depend only on G and are thus independent of the choice of parameters ε , δ , and θ (provided these fall within the legal range). We will again use the streamlined notation $\text{FP}(G)$ to denote the set of fixed point supports, as in section 3.

A few of our results express the relationship between $\text{FP}(G)$ and $\text{FP}(G')$ for a pair of related graphs G and G' (see, e.g., Rule 7), or have the form $\sigma \cup \{k\} \in \text{FP}(G)$ if and only if $\sigma \in \text{FP}(G)$ (e.g., Rule 8). Such relationships are also parameter independent, even when the sets $\text{FP}(G)$ and $\text{FP}(G')$ are themselves parameter dependent. In particular, for a different set of parameters, the fixed point supports $\text{FP}(G)$ and $\text{FP}(G')$ may change, but the given relationship between them stays the same. These statements should thus be understood as applying to CTLNs with fixed and matching parameters.

4.1 Uniform In-Degree and Parity. We begin with a very simple rule that follows directly from Theorem 1 (parity).

Rule 1 (parity). For any graph G , the total number of fixed points $|\text{FP}(G)|$ is odd.

Next, we digest Theorem 5 to obtain graph rules for some important special cases of uniform in-degree graphs: independent sets, cliques, and cycles. First, we repackage Theorem 5 as a rule to include it here for completeness.

Rule 2 (uniform in-degree). Let σ have uniform in-degree d . For $k \notin \sigma$, let $d_k \stackrel{\text{def}}{=} |\{i \in \sigma \mid i \rightarrow k\}|$ be the number of edges k receives from σ . Then

$$\sigma \in \text{FP}(G) \Leftrightarrow d_k \leq d \text{ for all } k \notin \sigma.$$

An *independent set* is a collection of nodes with no edges between them within the restricted subgraph; such a set is uniform in-degree with $d = 0$. Recall a *sink* is a node with no outgoing edges.

Rule 3 (independent sets). Suppose $\sigma \subseteq [n]$ is an independent set. Then

$$\sigma \in \text{FP}(G) \Leftrightarrow \sigma \text{ is a union of sinks.}$$

Furthermore, when $\sigma \in \text{FP}(G)$, the fixed point is stable if and only if $|\sigma| = 1$. In addition, $\text{idx}(\sigma) = (-1)^{|\sigma|-1}$.

To see the index formula in Rule 3, observe that a singleton has index 1, and by Lemma 3 (alternation), the indices must alternate their signs according to the size of the independent set.

A subset σ is a *clique* if every pair of nodes has a bidirectional edge between them in $G|_\sigma$, and thus it has uniform in-degree $d = |\sigma| - 1$. We say that a node $k \notin \sigma$ is a *target* of σ when $i \rightarrow k$ for all $i \in \sigma$. We say σ is *target-free* if there is no node $k \in [n] \setminus \sigma$ that k is a target of σ .

Rule 4 (cliques). Suppose $\sigma \subseteq [n]$ is a clique. Then

$$\sigma \in \text{FP}(G) \Leftrightarrow \sigma \text{ is target-free.}$$

Furthermore, when $\sigma \in \text{FP}(G)$, the fixed point is always stable and has $\text{idx}(\sigma) = +1$.

Recall that a *cycle* is a cyclically symmetric subgraph with uniform in-degree 1. Applying Theorem 5 yields the following result:




Figure 6: (A) A (proper) source in the graph. (B) A sink in the graph. (C) A projector onto the gray region. (D) A target of the gray region. Dashed lines indicate that all back edges are optional.

Rule 5 (cycles). Suppose $\sigma \subseteq [n]$ is a cycle. Then

$$\sigma \in \text{FP}(G) \Leftrightarrow \text{for all } k \notin \sigma, k \text{ receives at most 1 edge from } \sigma.$$

Furthermore, when $\sigma \in \text{FP}(G)$, the fixed point is always unstable, but $\text{idx}(\sigma) = +1$.

The index in Rules 4 and 5 follows from the fact that the sum of the indices of fixed points is always 1 (see Theorem 1), together with the observation that no proper subset of a clique or a cycle can support a fixed point. In the case of cliques, every proper subset is a clique with a target and thus does not survive. For a cycle, it is easy to see that any proper subset either contains a domination relationship or is an independent set that does not survive by Rule 3.

4.2 Adding a Single Node. Here we consider how fixed point supports are affected by the addition of a single node to a graph; specifically, we can fully characterize the effect on the fixed points when the added node is a source, sink, projector, or target (see Figure 6). A *source* is a node that has no incoming edges. We say a source is *proper* if it has at least one outgoing edge. A *sink* is a node that has no outgoing edges. Recall that a node k is a *projector onto* σ if $k \rightarrow i$ for all $i \in \sigma$; note that there may or may not be edges back from σ to k . Finally, recall that a node k is a *target of* σ if it receives an edge from every node in σ (again there may or may not be back edges from k to σ).

The proofs of the remaining rules will make use of the following technical results from the previous three sections:

- Theorem 2 (sign conditions) characterizing $\sigma \in \text{FP}(G)$ via the signs of the s_i^σ
- Corollary 2, which presents equivalent conditions for $\sigma \in \text{FP}(G)$ in terms of survival of σ in intermediate subgraphs $G|_\tau$ for $\tau \supseteq \sigma$
- Theorem 3 (simply-added) and Corollary 4, which describe how s_i^σ factors when σ contains a simply-added split




Figure 7: Three example graphs with their respective $\text{FP}(G)$.

- Theorem 4 (graphical domination) guaranteeing that if σ has inside-in or outside-in domination, then $\sigma \notin \text{FP}(G)$, while the presence of inside-out domination of a node $k \notin \sigma$ ensures the survival of a permitted motif σ to $\text{FP}(G|_{\sigma \cup k})$.

The next two rules show that no fixed point support can contain a node that is a proper source in the restricted subgraph, and that any node that is a proper source in all of G can be removed without any effect on $\text{FP}(G)$.

Rule 6 (sources in σ). Let G be a graph on n nodes and $\sigma \subseteq [n]$. If there exists a $k \in \sigma$ such that k is a proper source in $G|_{\sigma}$, then $\sigma \notin \text{FP}(G)$. More generally, if there exists an $\ell \in [n]$ such that σ contains a proper source in $G|_{\sigma \cup \{\ell\}}$, then $\sigma \notin \text{FP}(G)$.

Proof. Suppose there exists an $\ell \in [n]$ such that $k \in \sigma$ is a proper source in $G|_{\sigma \cup \{\ell\}}$ with $k \rightarrow \ell$. Then ℓ graphically dominates k since k has no inputs in $G|_{\sigma \cup \{\ell\}}$ and $k \rightarrow \ell$. Hence, $\sigma \notin \text{FP}(G|_{\sigma \cup \{\ell\}})$ by Theorem 4, and so $\sigma \notin \text{FP}(G)$ by Corollary 2. \square

Rule 7 (sources in G). Let G be a graph on n nodes and $k \in [n]$. If k is a proper source in G , then

$$\text{FP}(G) = \text{FP}(G|_{[n] \setminus \{k\}}).$$

Proof. Since k is a proper source in G , there exists $\ell \in [n]$ such that $k \rightarrow \ell$. Thus, for any $\sigma \subseteq [n]$ with $k \in \sigma$, σ contains a proper source in $G|_{\sigma \cup \{\ell\}}$, and so $\sigma \notin \text{FP}(G)$ by Rule 6. For $\sigma \subseteq [n] \setminus \{k\}$, $\sigma \in \text{FP}(G)$ implies $\sigma \in \text{FP}(G|_{[n] \setminus \{k\}})$, by Corollary 2. Hence, $\text{FP}(G) \subseteq \text{FP}(G|_{[n] \setminus \{k\}})$. To see the reverse inclusion, let $\sigma \in \text{FP}(G|_{[n] \setminus \{k\}})$. For any $i \in \sigma$, i inside-out dominates k with respect to σ , and so by Theorem 4, σ survives and thus $\sigma \in \text{FP}(G)$. Therefore, $\text{FP}(G) \supseteq \text{FP}(G|_{[n] \setminus \{k\}})$. \square

As an illustration of Rule 7, consider the graph in Figure 7B, which is obtained from the graph in panel A by adding a proper source, node 4. Observe that $\text{FP}(G)$ is identical for these two graphs, since the fixed points of G are preserved on removal of the proper source.

Rule 8 shows that the addition of a sink k creates a new fixed point support $\sigma \cup \{k\}$ precisely when σ is a surviving fixed point of the original graph.

Rule 8 (sinks). Let G be a graph on n nodes and $k \in [n]$. If k is a sink in G , then for any nonempty $\sigma \subseteq [n] \setminus \{k\}$,

$$\sigma \cup \{k\} \in \text{FP}(G) \Leftrightarrow \sigma \in \text{FP}(G),$$

and $\text{idx}(\sigma \cup \{k\}) = -\text{idx}(\sigma)$.

Proof. Since k is a sink in G , k is a nonprojector onto $[n] \setminus \{k\}$, and for all $i \in [n] \setminus \{k\}$, we have $s_i^{\sigma \cup \{k\}} = -\delta s_i^\sigma$ by Corollary 4. Thus, for all $i, j \in \sigma$, we have $\text{sgn } s_i^{\sigma \cup \{k\}} = \text{sgn } s_j^{\sigma \cup \{k\}}$ if and only if $\text{sgn } s_i^\sigma = \text{sgn } s_j^\sigma$. Similarly, for all $\ell \in [n] \setminus (\sigma \cup \{k\})$, $\text{sgn } s_\ell^{\sigma \cup \{k\}} = -\text{sgn } s_i^{\sigma \cup \{k\}}$ if and only if $\text{sgn } s_\ell^\sigma = -\text{sgn } s_i^\sigma$. Hence all the signs for all $i \in [n] \setminus \{k\}$ have the proper relationship for $\sigma \cup \{k\}$ to be a fixed point support precisely when σ is a fixed point support. Thus, the only remaining sign to check is that of $s_k^{\sigma \cup \{k\}} = s_k^\sigma$. Observe that $\sigma \cup \{k\} \in \text{FP}(G)$ precisely when, for all $i \in \sigma$,

$$\text{sgn } s_k^{\sigma \cup \{k\}} = \text{sgn } s_i^{\sigma \cup \{k\}} = \text{sgn}(-\delta s_i^\sigma) = -\text{sgn } s_i^\sigma,$$

in which case, $\text{sgn } s_k^\sigma = -\text{sgn } s_i^\sigma$ and so $\sigma \in \text{FP}(G)$. Thus, the result holds and $\text{idx}(\sigma \cup \{k\}) = -\text{idx}(\sigma)$, as desired. \square

As an example, consider the graph in Figure 7A, which consists of a clique $\{1, 2\}$ and a sink node 3. By Rule 8 (sinks), $\{1, 2, 3\} \in \text{FP}(G)$ if and only if $\{1, 2\} \in \text{FP}(G)$. Since $\{1, 2\}$ is a target-free clique, it survives as a fixed point by Rule 4, and thus $\{1, 2, 3\} \in \text{FP}(G)$ as well. In contrast, in Figure 7C, the addition of the sink node 4 eliminates the fixed point supports $\{1, 2\}$ and $\{1, 2, 3\}$ by Rule 2 (uniform in-degree). Thus, we know that $\{1, 2, 4\}, \{1, 2, 3, 4\} \notin \text{FP}(G)$ by Rule 8.

For our next rule, we consider when the added node k is a projector onto σ . In contrast to when k is a sink, $\sigma \cup \{k\}$ will be a fixed point support precisely when σ does not survive the addition of k .

Rule 9 (projectors). Let G be a graph on n nodes and $k \in [n]$. If k is a projector onto $[n] \setminus \{k\}$, then for any nonempty $\sigma \subseteq [n] \setminus \{k\}$,

$$\sigma \cup \{k\} \in \text{FP}(G) \Leftrightarrow \sigma \notin \text{FP}(G) \text{ and } \sigma \in \text{FP}(G|_{[n] \setminus \{k\}}),$$

and $\text{idx}(\sigma \cup \{k\}) = \text{idx}(\sigma)$.

Proof. Since k is a projector onto $[n] \setminus \{k\}$, we have that $s_i^{\sigma \cup \{k\}} = \varepsilon s_i^\sigma$ for all $i \in [n] \setminus \{k\}$ (see Corollary 4). By the same logic as in the proof of Rule 8, all the signs for $i \in [n] \setminus \{k\}$ have the proper relationship for $\sigma \cup \{k\}$ to be a fixed point support precisely when σ is a fixed point support in $G|_{[n] \setminus \{k\}}$.

Thus, the only remaining sign to check is that of $s_k^{\sigma \cup \{k\}} = s_k^\sigma$. Observe that $\sigma \cup \{k\} \in \text{FP}(G)$ precisely when $\sigma \in \text{FP}(G|_{[n] \setminus \{k\}})$, and for all $i \in \sigma$,

$$\text{sgn } s_k^{\sigma \cup \{k\}} = \text{sgn } s_i^{\sigma \cup \{k\}} = \text{sgn } s_i^\sigma.$$

In this case, $\text{sgn } s_k^\sigma = \text{sgn } s_i^\sigma$ and so $\sigma \notin \text{FP}(G)$. Finally, $\text{idx}(\sigma \cup \{k\}) = \text{idx}(\sigma)$. \square

The next rule shows that if σ has a target, it can never support a fixed point.

Rule 10 (targets). Let G be a graph on n nodes and $\sigma \subseteq [n]$. If there exists $k \in [n] \setminus \sigma$ such that k is a target of σ , then $\sigma \notin \text{FP}(G)$.

Proof. For every $i \in \sigma$, k outside-in dominates i with respect to σ , and so $\sigma \notin \text{FP}(G)$. \square

As an illustration of Rule 10, notice that in Figure 7C, node 4 is a target of $\{1, 2\}$, and thus $\{1, 2\} \notin \text{FP}(G)$.

The last two rules characterize the collection of fixed point supports when a graph contains a node that is either fully bidirectionally connected to all other nodes or, at the other extreme, is isolated and has no connections to other nodes.

Rule 11 (bidirectionally connected nodes). Let G be a graph on n nodes and $k \in [n]$. If $k \leftrightarrow i$ for all $i \in [n] \setminus \{k\}$, then

$$\text{FP}(G) = \{\sigma \cup \{k\} \mid \sigma \in \text{FP}(G|_{[n] \setminus \{k\}})\}.$$

Thus, $|\text{FP}(G)| = |\text{FP}(G|_{[n] \setminus \{k\}})|$.

Proof. Observe that k is a target of $[n] \setminus \{k\}$, and so for any σ with $k \notin \sigma$, $\sigma \notin \text{FP}(G)$ by Rule 10. Since k is also a projector onto $[n] \setminus \{k\}$, Rule 9 applies, and so $\sigma \cup \{k\} \in \text{FP}(G)$ if and only if $\sigma \in \text{FP}(G|_{[n] \setminus \{k\}})$ since $\sigma \notin \text{FP}(G)$ for all such σ . \square

Recall that a node is *isolated* if it has no incoming and no outgoing edges.

Rule 12 (isolated nodes). Let G be a graph on n nodes and $k \in [n]$. If k is an isolated node in G , then

$$\text{FP}(G) = \text{FP}(G|_{[n] \setminus \{k\}}) \cup \{\sigma \cup \{k\} \mid \sigma \in \text{FP}(G|_{[n] \setminus \{k\}})\} \cup \{k\}.$$

Thus, $|\text{FP}(G)| = 2|\text{FP}(G|_{[n] \setminus \{k\}})| + 1$.

Proof. Let $\sigma \in \text{FP}(G|_{[n] \setminus \{k\}})$. For any $i \in \sigma$, we have i inside-out dominates k with respect to σ , and so $\sigma \in \text{FP}(G)$. Since k is isolated, it is also a sink, and thus by Rule 8, $\sigma \cup \{k\} \in \text{FP}(G)$ for all $\sigma \in \text{FP}(G|_{[n] \setminus \{k\}})$. Finally, $\{k\} \in \text{FP}(G)$ by Rule 3 since it is trivially an independent set that is a sink. \square

4.3 Parameter Independence in $FP(G)$. All the CTLN graph rules presented thus far, including graphical domination, are parameter independent. Thus, any arguments relying on these results are guaranteed to depend solely on the graph structure. In light of this, it is natural to ask to what extent $FP(G)$ is parameter independent. The proof of Theorem 6 shows that graph rules are sufficient to characterize the full set $FP(G)$ for all graphs on $n \leq 4$ nodes, and thus $FP(G)$ is parameter independent for all such graphs:

Theorem 6. *Let G be a graph on $n \leq 4$ nodes. Then $FP(G)$ is constant across all values of ε and δ in the legal range. In other words, $FP(G)$ is parameter independent.*

Proof. Figure 18 in appendix section A.1 gives all directed graphs of size $n \leq 3$ together with the s_i^σ values, for $\sigma = [n]$. In particular, it can be seen that the signs of the s_i^σ for these small graphs are all parameter independent. Applying Theorem 2 (sign conditions), we see that there are nine permitted motifs of size $|\sigma| \leq 3$. Equivalently, all nine permitted motifs can be identified by graph rules, and all the other motifs can be shown to be forbidden by inside-in graphical domination. The survival rules for the permitted motifs are also parameter independent. Seven of them are uniform in-degree, and thus their survival is parameter independent by Rule 2, while for the remaining two, either inside-out or outside-in graphical domination applies to every possible embedding of the motifs into larger graphs. It follows that all fixed point supports that are proper subsets of $[n]$ for $n \leq 4$ are parameter independent. Moreover, by Rule 1 (parity), we can always determine whether the full support $[n]$ is in $FP(G)$ from knowledge of which proper subsets are in $FP(G)$. We conclude that for $n \leq 4$, $FP(G)$ is parameter independent. \square

It turns out that among permitted motifs of size $|\sigma| = 4$, a few have parameter-dependent survival rules when embedded in graphs of size $n \geq 5$. Figure 8 shows all three permitted motifs of size 4 whose survival is parameter dependent, together with the embeddings where this dependence occurs. (For all other embeddings, the survival rules are parameter independent and can be derived from graphical domination.) In fact, the precise dependence on ε and δ can be easily computed (see example 4) and is also shown in Figure 8. In appendix section A.2, we give some example graphs on $n = 5$ that have parameter-dependent $FP(G)$ as a result of having one or more of these motifs embedded in a parameter-dependent manner.

Example 4. For the three graphs in Figure 8, we will use Theorem 2 (sign conditions) to work out the parameter dependence of the survival of these motifs. In each case, let $\sigma = \{1, 2, 3, 4\}$ and let node 5 be the one in red.




Figure 8: Parameter-dependent survival. The three permitted motifs with $|\sigma| = 4$ that have parameter-dependent survival rules. In each panel, the graph in the gray-shaded region corresponds to σ and the red node is an external node that has an impact on the survival of σ . Note that all edges from the red node back to σ do not have an impact on the survival of σ and are thus omitted for clarity. In graph A, an equivalent embedding has edges from vertices 3 and 4 to the red node, by symmetry. In each case, the polynomial inequality gives the precise condition under which σ survives.

- A. Let G be the graph in Figure 8A. Computing $s_i^\sigma = \det((I - W_\sigma)_i; \theta \mathbf{1})$ for all $i \in \sigma$, we obtain

$$s_1^\sigma = \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2), \quad s_2^\sigma = s_3^\sigma = \varepsilon\theta(\varepsilon^2 + \varepsilon\delta + \delta^2),$$

$$s_4^\sigma = \varepsilon\theta(2\varepsilon^2 + 3\varepsilon\delta + 2\delta^2).$$

Since $\text{sgn } s_i^\sigma = \text{sgn } s_j^\sigma$ for all $i, j \in \sigma$, we see that σ is permitted by Theorem 2 (sign conditions). In addition, we have $s_5^\sigma = \varepsilon\theta(\varepsilon^3 + \varepsilon^2\delta - \delta^3)$. Since σ survives if and only if $\text{sgn } s_5^\sigma = -\text{sgn } s_i^\sigma$ for all $i \in \sigma$, we see that σ survives precisely when $\varepsilon^3 + \varepsilon^2\delta - \delta^3 < 0$.

- B. Let G be the graph in Figure 8B. Here we have

$$s_1^\sigma = \varepsilon\delta^2\theta, \quad s_2^\sigma = s_3^\sigma = \varepsilon\theta(\varepsilon^2 + \varepsilon\delta + \delta^2), \quad s_4^\sigma = \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2),$$

and since the signs of the s_i^σ all agree, we see that σ is a permitted motif. Then since $s_5^\sigma = \varepsilon\theta(\varepsilon^3 + \varepsilon^2\delta - \delta^3)$, we again see that σ survives precisely when $\varepsilon^3 + \varepsilon^2\delta - \delta^3 < 0$.

- C. Let G be the graph in Figure 8C. For this graph, we have

$$s_1^\sigma = s_2^\sigma = s_3^\sigma = -\varepsilon^2\delta\theta, \quad s_4^\sigma = -\varepsilon^2\theta(\varepsilon + 2\delta),$$

showing that σ is a permitted motif. In addition, $s_5^\sigma = -\varepsilon^2\theta(\varepsilon^2 + \varepsilon\delta - \delta^2)$, yielding the condition for survival of σ : $\varepsilon^2 + \varepsilon\delta - \delta^2 < 0$.

For certain classes of graphs, we can get parameter independence of $\text{FP}(G)$ through at least $n = 5$. A directed graph is called *oriented* if it has no bidirectional edges. (Note that all the graphs in Figure 8 have the bidirectional edge $2 \leftrightarrow 3$ and are thus not oriented.)

Theorem 7. *Let G be an oriented graph on $n \leq 5$ nodes. Then $FP(G)$ is parameter independent.*

Proof. In appendix section A.3, we examine all permitted motifs with $|\sigma| \leq 4$ that can arise as subgraphs of an oriented graph and show that the corresponding survival rules are all parameter independent. Combining this with Rule 1 (parity), we see that the full $FP(G)$ is parameter independent for oriented graphs of size $n \leq 5$. \square

We have already seen that parameter independence of $FP(G)$ cannot extend to all $n = 5$ graphs (see appendix section A.2 for some explicit examples). It is possible, however, that Theorem 7 can be extended to oriented graphs with $n > 5$.

5 Building Block Graph Rules for CTLNs

In sections 3 and 4, we have presented a variety of graph rules that enable us to determine when certain subgraphs are permitted or forbidden motifs, as well as their survival rules. In this section, we consider larger networks corresponding to composite graphs that are built from component subgraphs and address the question: If we know $FP(G_i)$ for each component G_i , what can we say about $FP(G)$ for the full network? Our main results are collected in section 5.1. These results can be thought of as graph rules for networks that can be decomposed into smaller building blocks whose structure is better understood. In the remaining subsections, we prove these results and provide additional details on the fixed point supports of composite graphs.

As in sections 3 and 4, we simplify notation and write $FP(G)$ instead of $FP(G, \varepsilon, \delta)$. Note that whenever such expressions occur in the same claim, such as $FP(G)$ and $FP(G_i)$, the same choices for ε and δ must be assumed in each instance. All the results, however, are parameter independent.

5.1 Composite Graphs.

Definition 8 (*composite graph*). *Given a set of graphs G_1, \dots, G_N , and a graph \widehat{G} on N nodes, the composite graph with components G_i and skeleton \widehat{G} is the graph G constructed by taking the union of all component graphs and adding edges between components according to the following rule: if $u \in G_i$ and $v \in G_j$, then $u \rightarrow v$ in G if and only if $i \rightarrow j$ in \widehat{G} . (See Figure 9.)*

Note that any graph G can trivially be thought of as a composite graph, with each component a single vertex. The results in this section are more interesting when at least one component in the graph is larger.

The first result says that if $G|_\sigma$ is a composite graph, for some $\sigma \subseteq [n]$, and any component σ_i is a forbidden motif, then σ is also a forbidden motif. In other words, one bad apple spoils the bunch. The proof is given in section 5.2.




Figure 9: (A) A skeleton graph \widehat{G} . (B) An arbitrary composite graph with skeleton \widehat{G} from panel A. Each node i in the skeleton is replaced with a component graph G_i whose connections to the rest of the graph are prescribed by the connections of node i in \widehat{G} . (C) An example composite graph with skeleton \widehat{G} from panel A.

Theorem 8. *Let $G|_\sigma$ be any composite graph with components $G|_{\sigma_1}, \dots, G|_{\sigma_N}$. If σ_i is a forbidden motif for any $i \in [N]$, then σ is also a forbidden motif.*

Next we show that if $G|_\sigma$ is a composite graph that has an unstable uniform in-degree component, this is sufficient to guarantee that σ is also unstable.

Proposition 1. *Let $G|_\sigma$ be any composite graph with components $G|_{\sigma_1}, \dots, G|_{\sigma_N}$. If σ_i is a uniform in-degree permitted motif that is unstable for any $i \in [N]$, then σ is unstable (or rather $I - W_\sigma$ is unstable).*

Proof. By definition of a composite graph, σ has a simply-added split where $\sigma \setminus \sigma_i$ is simply-added to σ_i . Since there exists a σ_i that is uniform in-degree, we can apply Lemma 5 to show that all the eigenvalues of $I - W_{\sigma_i}$ are inherited to $I - W_\sigma$ except for possibly the top eigenvalue (which is the row sum of $I - W_{\sigma_i}$). Since σ_i is unstable, $I - W_{\sigma_i}$ has at least one negative eigenvalue, and this cannot be the top eigenvalue, since the row sum is always positive. Thus, $I - W_\sigma$ inherits this negative eigenvalue, and so is unstable. \square

For the following results, G is always a composite graph on n vertices, with skeleton \widehat{G} and components G_1, \dots, G_N . For any $\sigma \subseteq [n]$, we define $\sigma_i \stackrel{\text{def}}{=} \sigma \cap \tau_i$, where τ_i is the set of vertices in component G_i . Note that $[n] = \cup_i \tau_i$, and $G_i = G|_{\tau_i}$.

Recall that a graph is *strongly forbidden* if it has inside-in graphical domination. The next theorem shows that when the skeleton is strongly forbidden, the full graph is as well:

Theorem 9. *Let G be a composite graph with skeleton \widehat{G} . If \widehat{G} is strongly forbidden, then G is strongly forbidden.*

A strongly forbidden skeleton is thus sufficient to guarantee that G is strongly forbidden even if all the components G_i are permitted. Graphical domination within the skeleton is essential to this result: the following




Figure 10: (A1) A strongly forbidden motif, where node 1 graphically dominates node 3 with respect to $\{1, 2, 3\}$. (A2) A composite graph whose skeleton is the graph from panel A1. This graph is a strongly forbidden motif since every node in component G_1 graphically dominates every node in G_3 . (B1) A forbidden motif that has no graphical domination. (B2) A composite graph whose skeleton is the graph from panel B1. This graph is a (uniform in-degree) permitted motif despite having a forbidden skeleton.

example shows that Theorem 9 cannot be extended to skeletons that are forbidden but not strongly forbidden.

Example 5. Consider the graph in panel A1 of Figure 10. This graph is strongly forbidden since node 1 graphically dominates node 3 with respect to the full support $\{1, 2, 3\}$. The graph in panel A2 of Figure 10 is a composite obtained from the skeleton in panel A1 by inserting 2-cliques at two of the nodes. Theorem 9 guarantees that the graph in panel A2 is strongly forbidden since it has a strongly forbidden skeleton. In particular, we see that every node in G_1 graphically dominates each node in G_3 .

In contrast, consider the graph in panel B1 of Figure 10, which has no graphical domination and thus is not strongly forbidden. It turns out that this graph is still a forbidden motif: it has the form $\sigma \cup \{k\}$ where $\sigma = \{1, 2, 3, 4\}$ and $k = 5$ is a sink. Since σ has uniform in-degree $d = 1$ and 5 receives two edges from σ , we see $\sigma \notin \text{FP}(G)$ by Rule 2 and thus $\sigma \cup \{5\} \notin \text{FP}(G)$ by Rule 8 (sinks). The composite graph in panel B2 of Figure 10 has the graph from panel B1 as its skeleton, again with 2-cliques inserted at two of the nodes. However, this graph is not forbidden, despite having a forbidden skeleton, since it has uniform in-degree 2 (see Rule 2). This example shows that Theorem 9 cannot be generalized from skeletons that are strongly forbidden to those that are merely forbidden.

Next we consider some special classes of composite graphs where we can guarantee that the full support is permitted precisely when all the components are permitted.

Definition 9. Let G be a composite graph with components G_1, \dots, G_N and skeleton \widehat{G} . If \widehat{G} is an independent set, then we say that G is the disjoint union of the components G_i . If \widehat{G} is a clique, then we say G is the clique union of the G_i . Finally, if \widehat{G} is a cycle, then G is a cyclic union of its components. (See Figure 11.)

Theorem 10. Let $G|_\sigma$ be either a disjoint union, a clique union, or a cyclic union with nonempty components $G|_{\sigma_1}, \dots, G|_{\sigma_N}$. Then σ is a permitted motif if and only




Figure 11: (A) A disjoint union. (B) A clique union. (C) A cyclic union.

if σ_i is a permitted motif for every $i \in [N]$. Moreover, when σ is permitted, the index factors as:

$$idx(\sigma) = \begin{cases} (-1)^{N-1} \prod_{i=1}^N idx(\sigma_i), & \text{for disjoint unions,} \\ \prod_{i=1}^N idx(\sigma_i), & \text{for clique or cyclic unions.} \end{cases}$$

Note that the expressions for the indices, $idx(\sigma)$, are identical or nearly identical between clique unions, cyclic unions, and disjoint unions. A simple rule of thumb is that the index of σ in these unions is given by the product of the indices of all components σ_i with the index of the skeleton. In the case of clique unions and cyclic unions, the skeletons have a unique (full support) fixed point, whose index is thus +1. For disjoint unions, the skeleton is an independent set of size N with index $(-1)^{N-1}$ (see Rule 3).

Disjoint unions, clique unions, and cyclic unions are all examples of composite graphs whose skeletons are permitted motifs. It is thus natural to ask whether Theorem 10 holds for other composite graphs with permitted skeletons. In fact, we conjecture that this theorem holds more generally.

Conjecture 1. Let G be a composite graph whose skeleton \widehat{G} is a permitted motif. Then G is a permitted motif if and only if every component G_i is permitted.

Note that the forward direction is a direct consequence of Theorem 8. It is also worth noting that even if true, this conjecture does not cover all permitted motifs that are composite graphs. For example, the graph in panel B2 of Figure 10 is a permitted motif, but its skeleton is forbidden.

The proof of Theorem 10 is given in section 5.3. There we provide separate proofs for the disjoint, clique, and cyclic unions and also work out the full sets $FP(G)$ in the case where all of G has one of these special composite structures. Furthermore, these results imply that $FP(G)$ is parameter independent whenever all the $FP(G_i)$ are parameter independent (see

Corollary 5). In section 5.4, we consider survival rules for disjoint unions and clique unions that are embedded in larger graphs.

But first, in section 5.2, we present some lemmas that help us to prove the theorems on composite graphs stated above. There we also prove Theorems 8 and 9.

5.2 Some Lemmas for Composite Graphs.

Lemma 6. *Let G be a composite graph, and consider $\sigma \subseteq [n]$. For any component G_i such that $\sigma_i = \sigma \cap \tau_i \neq \emptyset$,*

$$\text{sgn } s_j^\sigma = \text{sgn } s_k^\sigma \Leftrightarrow \text{sgn } s_j^{\sigma_i} = \text{sgn } s_k^{\sigma_i}, \text{ for all } j, k \in \tau_i.$$

Proof. Observe that $[n] \setminus \tau_i$ is simply-added to τ_i for each $i \in [N]$, by the nature of the composite graph construction, and thus $\sigma \setminus \sigma_i$ is simply-added to τ_i . By Theorem 3, $s_j^\sigma = \alpha s_j^{\sigma_i}$, where $\alpha = \frac{1}{\theta} s_j^{\sigma \setminus \sigma_i}$ is identical for all $j \in \tau_i$. \square

With this lemma, we can immediately prove Theorem 8.

Proof of Theorem 8. Suppose $G|_\sigma$ is a composite graph with components $G|_{\sigma_1}, \dots, G|_{\sigma_N}$, and one of the σ_i is a forbidden motif. Then, by Theorem 2 (sign conditions), there exists $j, k \in \sigma_i$ such that $\text{sgn } s_j^{\sigma_i} \neq \text{sgn } s_k^{\sigma_i}$, and therefore by Lemma 6, we have $\text{sgn } s_j^\sigma \neq \text{sgn } s_k^\sigma$. Thus, σ is a forbidden motif. \square

Lemma 7. *If $\sigma \in \text{FP}(G)$, then $\sigma_i \in \text{FP}(G_i)$ for each $i \in [N]$ such that $\sigma_i \neq \emptyset$.*

Proof. If $\sigma \in \text{FP}(G)$, then for any $j, k \in \sigma_i$ and $\ell \in \tau_i \setminus \sigma_i$, we have $\text{sgn } s_j^\sigma = \text{sgn } s_k^\sigma = -\text{sgn } s_\ell^\sigma$ by Theorem 2 (sign conditions). By Lemma 6, we infer that $\text{sgn } s_j^{\sigma_i} = \text{sgn } s_k^{\sigma_i} = -\text{sgn } s_\ell^{\sigma_i}$, and so σ_i satisfies the sign conditions within G_i . Thus, $\sigma_i \in \text{FP}(G_i)$. \square

As an immediate consequence of Lemma 7, we see that if any component σ_i of σ is forbidden (or non-surviving) in G_i , then σ is forbidden (or non-surviving) in G . The next lemma shows that if σ is a permitted motif, the survival of each of its components in G is sufficient to guarantee the survival of σ .

Lemma 8. *If $\sigma \in \text{FP}(G|_\sigma)$ and $\sigma_i \in \text{FP}(G_i)$ for each $i \in [N]$, then $\sigma \in \text{FP}(G)$.*

Proof. First, note that $\sigma_i \in \text{FP}(G_i)$ implies $\sigma_i \neq \emptyset$ for each $i \in [N]$, and $\text{sgn } s_j^{\sigma_i} = \text{sgn } s_k^{\sigma_i} = -\text{sgn } s_\ell^{\sigma_i}$ for all $j, k \in \sigma_i$ and $\ell \in \tau_i \setminus \sigma_i$. On the other hand, $\sigma \in \text{FP}(G|_\sigma)$ implies $\text{sgn } s_j^\sigma = \text{sgn } s_k^\sigma$ for all $j, k \in \sigma$. Now applying Lemma 6, we see that $\text{sgn } s_\ell^\sigma = -\text{sgn } s_j^\sigma$ for $\ell \in [n] \setminus \sigma$ and $j \in \sigma$. It follows that the sign conditions are satisfied, and so $\sigma \in \text{FP}(G)$. \square

The next lemma tells us that graphical domination in the skeleton graph \widehat{G} gets inherited to the composite graph G .

Lemma 9. *Let G be a composite graph with skeleton \widehat{G} . If \widehat{k} graphically dominates \widehat{j} with respect to $\widehat{\sigma}$ in \widehat{G} , then for any $k \in G_{\widehat{k}}$ and any $j \in G_{\widehat{j}}$, we have that k graphically dominates j with respect to $\sigma = \bigcup_{i \in \widehat{\sigma}} G_i$.*

Proof. Let $[n]$ denote the vertices of G and $[N]$ the vertices of \widehat{G} . Consider $\widehat{\sigma} \subseteq [N]$, and suppose there exists $\widehat{j}, \widehat{k} \in [N]$ such that \widehat{k} graphically dominates \widehat{j} with respect to $\widehat{\sigma}$. By definition of graphical domination, this means that (1) for any $\widehat{i} \in \widehat{\sigma}$, $\widehat{i} \rightarrow \widehat{j}$ implies $\widehat{i} \rightarrow \widehat{k}$; (2) if $\widehat{j} \in \widehat{\sigma}$, then $\widehat{j} \rightarrow \widehat{k}$; and (3) if $\widehat{k} \in \widehat{\sigma}$, then $\widehat{k} \not\rightarrow \widehat{j}$. We will show that in G , any $j \in G_{\widehat{j}}$ and $k \in G_{\widehat{k}}$ satisfy these same properties, and so k graphically dominates j with respect to σ .

By definition of the composite graph, $\widehat{j} \rightarrow \widehat{k}$ and $\widehat{k} \not\rightarrow \widehat{j}$ in the skeleton immediately imply $j \rightarrow k$ and $k \not\rightarrow j$ in $G|_{\sigma}$, and so satisfying conditions 2 and 3 in the skeleton guarantees they are satisfied in G . Now we turn to condition 1. Consider $i \in \sigma$ such that $i \rightarrow j$ in $G|_{\sigma}$. If i and j are both in component $\sigma_{\widehat{j}}$, then $j \in \sigma$ and so $\widehat{j} \in \widehat{\sigma}$. Thus, condition 2 holds and $\widehat{j} \rightarrow \widehat{k}$ in $\widehat{G}|_{\widehat{\sigma}}$, which implies that $i \rightarrow k$ in $G|_{\sigma}$. If they are in different components (i.e., $i \in \sigma_{\widehat{i}}$ for some $\widehat{i} \neq \widehat{j}$), then by definition of a composite graph, $i \rightarrow j$ implies $\widehat{i} \rightarrow \widehat{j}$ in $\widehat{G}|_{\widehat{\sigma}}$, and so by condition 1, we must have $\widehat{i} \rightarrow \widehat{k}$ and thus $i \rightarrow k$ in $G|_{\sigma}$. Thus, for any $i \in \sigma$, $i \rightarrow j$ implies $i \rightarrow k$ and condition 1 is satisfied. Hence, k graphically dominates j with respect to σ , and so $G|_{\sigma}$ is strongly forbidden. \square

As an immediate consequence, we obtain Theorem 9, since inside-in graphical domination within any graph implies that it is strongly forbidden.

5.3 Proof of Theorem 10 (Disjoint, Clique, and Cyclic Unions). To prove Theorem 10, we separately consider composite graphs that are disjoint unions, clique unions, and cyclic unions. We characterize the full FP(G) of these graphs. In the next section, we provide partial survival rules when these motifs are embedded in larger graphs.

As before, in this section we assume G is a composite graph on n vertices, with skeleton \widehat{G} and components G_1, \dots, G_N . For any $\sigma \subseteq [n]$, we denote $\sigma_i = \sigma \cap \tau_i$, where τ_i is the set of vertices in G_i . Note that $[n] = \bigcup_i \tau_i$, and $G_i = G|_{\tau_i}$.

Theorem 11 (disjoint union). *Let G be a disjoint union of components G_1, \dots, G_N . For any nonempty $\sigma \subseteq [n]$,*

$$\sigma \in \text{FP}(G) \iff \sigma_i \in \text{FP}(G_i) \cup \emptyset \text{ for all } i \in [N].$$

Moreover, for any $\sigma \in \text{FP}(G)$, the index is given by

$$\text{id}x(\sigma) = (-1)^{|\widehat{\sigma}|-1} \prod_{i \in \widehat{\sigma}} \text{id}x(\sigma_i), \quad \text{where } \widehat{\sigma} \stackrel{\text{def}}{=} \{i \in [N] \mid \sigma_i \neq \emptyset\}.$$

In particular, Theorem 11 tells us that

$$\text{FP}(G) = \{\sigma \subseteq [n] \mid \sigma_i \in \text{FP}(G_i) \cup \emptyset \text{ for each } i \in [N]\},$$

and so,

$$|\text{FP}(G)| = \prod_{i=1}^N (|\text{FP}(G_i)| + 1) - 1.$$

For example, if $|\text{FP}(G_i)| = 1$ for each $i \in [N]$, as in the case of the independent set (where every component is a single node), then $|\text{FP}(G)| = 2^N - 1$.

Proof. (\Rightarrow) This follows immediately from Lemma 7.

(\Leftarrow) To simplify notation, we fix $\theta = 1$.⁵ Suppose σ has the property that $\sigma_i \in \text{FP}(G_i)$ for each $i \in \widehat{\sigma}$, where $\widehat{\sigma} \stackrel{\text{def}}{=} \{i \in [N] \mid \sigma_i \neq \emptyset\}$. Consider $j \in \sigma_i$, and observe that for any $k \neq i$, σ_k is simply-added to σ_i . Using Theorem 3, we thus obtain $s_j^{\sigma_i \cup \sigma_k} = s_j^{\sigma_i} s_j^{\sigma_k}$. Moreover, for any other $\ell \in \widehat{\sigma} \setminus \{i, k\}$, we also have that σ_ℓ is simply-added to $\sigma_i \cup \sigma_k$, and so

$$s_j^{\sigma_i \cup \sigma_k \cup \sigma_\ell} = s_j^{\sigma_i} s_j^{\sigma_k} s_j^{\sigma_\ell}.$$

Continuing in this fashion yields

$$s_j^\sigma = \prod_{\ell \in \widehat{\sigma}} s_j^{\sigma_\ell}.$$

Furthermore, the above formula also holds for $j \notin \sigma$. This is because $j \in \tau_k$ for some k , and all other (nonempty) σ_i are simply-added to τ_k , not just to σ_k .

Note that for each $\sigma_i \in \text{FP}(G_i)$, σ_i survives as a fixed point in $\text{FP}(G)$ by inside-out domination. Therefore, for any $j \in [n]$, $\text{sgn} s_j^{\sigma_i} = \text{id}x(\sigma_i)$ if and only if $j \in \sigma_i$. For $j \in \sigma$, we have $j \in \sigma_i$ for exactly one $i \in \widehat{\sigma}$, and thus

$$\text{sgn} s_j^\sigma = \text{sgn} s_j^{\sigma_i} \prod_{\ell \in \widehat{\sigma} \setminus i} \text{sgn} s_j^{\sigma_\ell} = \text{id}x(\sigma_i) \prod_{\ell \in \widehat{\sigma} \setminus i} (-\text{id}x(\sigma_\ell)) = (-1)^{|\widehat{\sigma}|-1} \prod_{\ell \in \widehat{\sigma}} \text{id}x(\sigma_\ell).$$

⁵Note that $\text{FP}(G)$ never depends on the value of θ , provided $\theta > 0$, so this choice cannot affect the results.

Since the sign of s_j^σ is the same for all $j \in \sigma$, we see σ is permitted and the index $\text{idx}(\sigma)$ must match the sign of s_j^σ (by Theorem 2). On the other hand, for $j \notin \sigma$,

$$\text{sgn } s_j^\sigma = \prod_{\ell \in \hat{\sigma}} \text{sgn } s_j^{\sigma_\ell} = \prod_{\ell \in \hat{\sigma}} (-\text{idx}(\sigma_\ell)) = (-1)^{|\hat{\sigma}|} \prod_{\ell \in \hat{\sigma}} \text{idx}(\sigma_\ell) = -\text{idx}(\sigma).$$

By Theorem 2, we conclude that $\sigma \in \text{FP}(G)$ with $\text{idx}(\sigma) = (-1)^{|\hat{\sigma}|-1} \prod_{\ell \in \hat{\sigma}} \text{idx}(\sigma_\ell)$. \square

The next theorem characterizes the full $\text{FP}(G)$ whenever G is a clique union.

Theorem 12 (*clique union*). *Let G be a graph on n vertices that is a clique union of N subgraphs G_1, \dots, G_N . For any $\sigma \subseteq [n]$,*

$$\sigma \in \text{FP}(G) \iff \sigma_i \in \text{FP}(G_i) \text{ for all } i \in [N].$$

Moreover, if $\sigma \in \text{FP}(G)$, then $\text{idx}(\sigma) = \prod_{i=1}^N \text{idx}(\sigma_i)$.

Note that the theorem implies that the total number of fixed points satisfies

$$|\text{FP}(G)| = \prod_{i=1}^N |\text{FP}(G_i)|.$$

In particular, if each component G_i has a unique fixed point, then G has a unique fixed point.

Proof. (\Rightarrow) Suppose $\sigma \in \text{FP}(G)$. Lemma 7 guarantees that for any $i \in [N]$, $\sigma_i \in \text{FP}(G_i)$ or $\sigma_i = \emptyset$. If we assume there exists an $i \in [N]$ such that $\sigma_i = \emptyset$, then any $k \in \tau_i$ is a target of σ , and thus by Rule 10, we cannot have $\sigma \in \text{FP}(G)$, a contradiction. We thus conclude that $\sigma_i \neq \emptyset$, and so $\sigma_i \in \text{FP}(G_i)$, for all $i \in [N]$.

(\Leftarrow) To simplify notation, we fix $\theta = 1$. Suppose $\sigma_i \in \text{FP}(G_i)$ for all $i \in [N]$. In particular, each σ_i is nonempty. As with disjoint unions, clique unions have the property that each σ_ℓ is simply-added to any union of the other components. We thus have the same formula as in the proof of Theorem 11, where for any $j \in [n]$ and $\sigma \subseteq [n]$ the value s_j^σ factors as

$$s_j^\sigma = \prod_{\ell \in [N]} s_j^{\sigma_\ell}.$$

Note that here, the product is necessarily over all $\ell \in [N]$, since each σ_ℓ is nonempty. Moreover, for any $\ell \in [N]$ such that $j \notin \tau_\ell$, j is a target of σ_ℓ , and

so $\sigma_\ell \notin \text{FP}(G|_{\sigma_i \cup j})$. This implies, by Theorem 2, that $\text{sgn } s_j^{\sigma_\ell} = \text{idx}(\sigma_\ell)$ for all $j \notin \tau_\ell$.

To show that $\sigma \in \text{FP}(G)$, we must show that $\text{sgn } s_j^\sigma = \text{idx}(\sigma)$ for each $j \in \sigma$, and has the opposite sign if $j \notin \sigma$ (see Theorem 2). Consider $j \in \sigma$, and observe that $j \in \sigma_i$ for some i . Clearly, $\text{sgn } s_j^{\sigma_i} = \text{idx}(\sigma_i)$. By the argument above, we also have $\text{sgn } s_j^{\sigma_\ell} = \text{idx}(\sigma_\ell)$ for each $\ell \in [N] \setminus i$. Therefore, using the above product formula for s_j^σ , we have $\text{sgn } s_j^\sigma = \prod_{\ell=1}^N \text{idx}(\sigma_\ell)$. Since the sign matches for all $j \in \sigma$, we have that σ is permitted and also that $\text{idx}(\sigma) = \prod_{\ell=1}^N \text{idx}(\sigma_\ell)$, as desired.

On the other hand, for any $j \notin \sigma$, there exists an i such that $j \in \tau_i \setminus \sigma_i$. Since $\sigma_i \in \text{FP}(G_i)$, $\text{sgn } s_j^{\sigma_i} = -\text{idx}(\sigma_i)$, while for all other $\ell \in [N] \setminus i$, we have $j \notin \tau_\ell$ and so $\text{sgn } s_j^{\sigma_\ell} = \text{idx}(\sigma_\ell)$. Thus,

$$\text{sgn } s_j^\sigma = -\text{idx}(\sigma_i) \prod_{\ell \in [N] \setminus i} \text{idx}(\sigma_\ell) = - \prod_{\ell=1}^N \text{idx}(\sigma_\ell) = -\text{idx}(\sigma).$$

By Theorem 2, we conclude that $\sigma \in \text{FP}(G)$. □

We conclude this section with Theorem 13 and its proof, characterizing $\text{FP}(G)$ when G is a cyclic union.

Theorem 13 (*cyclic union*). *Let G be a cyclic union of components G_1, \dots, G_N . Then*

$$\sigma \in \text{FP}(G) \Leftrightarrow \sigma_i \in \text{FP}(G_i) \text{ for all } i \in [N].$$

Moreover, if $\sigma \in \text{FP}(G)$, then $\text{idx}(\sigma) = \prod_{i=1}^N \text{idx}(\sigma_i)$.

Note that, just as with the clique union, the theorem implies that for cyclic unions,

$$|\text{FP}(G)| = \prod_{i=1}^N |\text{FP}(G_i)|.$$

To prove Theorem 13, we will need the following lemmas:

Lemma 10. *Let G be a cyclic union of components G_1, \dots, G_N . If $\sigma \in \text{FP}(G)$, then $\sigma_i \neq \emptyset$ for all $i \in [N]$.*

Proof. Let $\sigma \in \text{FP}(G)$ and $\hat{\sigma} = \{i \in [N] \mid \sigma_i \neq \emptyset\}$. Suppose for the sake of contradiction that $\hat{\sigma} \neq [N]$. Then the skeleton $\widehat{G}|_{\hat{\sigma}}$ of σ is a proper subset of a cycle, and so the skeleton must either contain a proper source or be an independent set. If $\widehat{G}|_{\hat{\sigma}}$ contains a proper source, then $\hat{\sigma}$ has inside-in graphical domination, and so by Lemma 9, that same domination lifts to

inside-in graphical domination in σ . If $\hat{\sigma}$ is an independent set, then there is outside-in graphical domination of $\hat{\sigma}$ by any external node that receives an edge from $\hat{\sigma}$. Thus, again by Lemma 9, that domination relationship lifts to outside-in graphical domination of σ . But then $\sigma \notin \text{FP}(G)$ by Theorem 4, yielding a contradiction. Thus, we must have $\hat{\sigma} = [N]$, that is, $\sigma_i \neq \emptyset$ for all $i \in [N]$. \square

The next lemma is an immediate consequence of Lemmas 7 and 8.

Lemma 11. *Let $\sigma \subseteq [n]$ have nontrivial overlap with each component of G , so that $\sigma_i \neq \emptyset$ for each $i \in [N]$. Then*

$$\sigma \in \text{FP}(G) \Leftrightarrow \sigma \in \text{FP}(G|_{\sigma}) \text{ and } \sigma_i \in \text{FP}(G_i) \text{ for each } i \in [N].$$

We are now ready to prove the cyclic union theorem.

Proof of Theorem 13 (cyclic union). We prove this by complete induction on n , the total number of vertices of G , and for a fixed number of components N . The base case is $n = N$, so that G is an n -cycle. In this case, there is a unique fixed point of full support (by the remarks following Rule 5), and every component is a single vertex with index 1. The result trivially holds.

Now assume the inductive hypothesis that the theorem statement holds for all cyclic unions with N components and m vertices, where $N \leq m < n$. By Lemmas 10 and 11,

$$\sigma \in \text{FP}(G) \Leftrightarrow \sigma \in \text{FP}(G|_{\sigma}) \text{ and } \sigma_i \in \text{FP}(G_i) \text{ for all } i \in [N].$$

So to prove the first theorem statement for a graph of n vertices, it suffices to show that if $\sigma_i \in \text{FP}(G_i)$ for all $i \in [N]$, then $\sigma \in \text{FP}(G|_{\sigma})$.

Suppose $|\sigma| < n$, and suppose $\sigma_i \in \text{FP}(G_i)$ for all $i \in [N]$. Recall that if $\sigma_i \in \text{FP}(G_i)$, then $\sigma_i \in \text{FP}(G|_{\sigma_i})$. Applying the inductive hypothesis to $G|_{\sigma}$, we see that this in turn implies $\sigma \in \text{FP}(G|_{\sigma})$, as desired. Moreover, the index formula holds because $\text{idx}(\sigma)$ is the same as what it was in the smaller graph $G|_{\sigma}$ and is thus given immediately by the inductive hypothesis.

Now assume $|\sigma| = n$. This means $G|_{\sigma} = G$, and $\sigma_i = \tau_i$ for each $i \in [N]$. By Lemma 11, $\sigma \in \text{FP}(G)$ implies $\sigma_i \in \text{FP}(G_i)$ for all $i \in [N]$. So what remains is to show the converse direction and that the index formula holds. Suppose $\tau_i \in \text{FP}(G_i)$ for each $i \in [N]$. If $\sigma = [n] \notin \text{FP}(G)$, then $|\text{FP}(G)| = \prod_{i \in [N]} |\text{FP}(G_i)| - 1$, since all the smaller elements of $\text{FP}(G)$ are indeed given by picking a fixed point support from each component graph G_i . By parity, each $|\text{FP}(G_i)|$ is odd, and hence the product of these terms is odd. It follows that $|\text{FP}(G)|$ is even, contradicting parity for G . We conclude, then, that we must have $\sigma \in \text{FP}(G)$.

Finally, we show that the index formula holds for $\sigma = [n]$ (by assumption, it holds for $|\sigma| < n$). Using Theorem 1 (the index theorem) for $\text{FP}(G)$,




Figure 12: The two skeletons \widehat{G} considered in Lemma 12.

we compute

$$\begin{aligned}
 1 &= \sum_{\sigma \in \text{FP}(G)} \text{idx}(\sigma) = \sum_{\sigma \in \text{FP}(G) \setminus [n]} \text{idx}(\sigma) + \text{idx}([n]) \\
 &= \sum_{\sigma \in \text{FP}(G) \setminus [n]} \left(\prod_{i \in [N]} \text{idx}(\sigma_i) \right) + \text{idx}([n]) \\
 &= \prod_{i \in [N]} \left(\sum_{\omega \in \text{FP}(G_i)} \text{idx}(\omega) \right) - \prod_{i \in [N]} \text{idx}(\tau_i) + \text{idx}([n]) \\
 &= 1 - \prod_{i \in [N]} \text{idx}(\tau_i) + \text{idx}([n]).
 \end{aligned}$$

Note that in the last equality, we have used the index theorem to see that the indices sum to 1 for each $\text{FP}(G_i)$. It follows that $\text{idx}([n]) = \prod_{i \in [N]} \text{idx}(\tau_i)$, as desired. \square

It is worth thinking about what properties of cycles made the induction argument in theorem 13 work. Namely, we needed \widehat{G} to be a permitted motif and for every proper subset of its vertices to be either strongly forbidden or guaranteed not to survive as a fixed point support due to outside-in graphical domination, as in Lemma 10. In other words, $\text{FP}(\widehat{G})$ must have a unique fixed point that has full support, with all proper subsets excluded via graphical domination. Note that cliques also have this property, so Theorem 12 could have been proven with a similar inductive argument.

The following lemma provides an example of how the inductive argument can generalize beyond clique unions and cyclic unions.

Lemma 12. *Let G be a composite graph with one of the skeletons \widehat{G} depicted in Figure 12, with $N = 4$ or $N = 5$. Then*

$$\sigma \in FP(G) \Leftrightarrow \sigma_i \in FP(G_i) \text{ for all } i \in [N].$$

Moreover, if $\sigma \in FP(G)$, then $idx(\sigma) = \prod_{i=1}^N idx(\sigma_i)$.

Proof sketch. First, consider the $N = 4$ skeleton in Figure 12A. In this case, \widehat{G} has uniform in-degree 2 and is thus a permitted motif. It is straightforward to check that all proper subsets of \widehat{G} are either strongly forbidden or do not survive as fixed point supports due to outside-in graphical domination. Thus, an analogue of Lemma 10 holds, and Lemma 11 applies. The induction proof of Theorem 13 can therefore be easily adapted to this case, yielding the desired result.

For the $N = 5$ skeleton (see Figure 12B), we can again check that all proper subsets of \widehat{G} are either strongly forbidden or do not survive as fixed point supports due to outside-in graphical domination. By Rule 1 (parity), the full graph must be a permitted motif. The remaining arguments then follow exactly as in the previous case. \square

We are now ready to prove Theorem 10. Putting together Theorems 11 to 13, we immediately obtain the proof.

Proof of Theorem 10. To see the first statement, apply Theorems 11 to 13 for $G = G|_\sigma$ a disjoint union, a clique union, and a cyclic union, respectively. Since each σ_i is nonempty, it follows in every case that $\sigma \in FP(G|_\sigma)$ if and only if $\sigma_i \in FP(G|_{\sigma_i})$ for all $i \in [N]$. Hence, σ is a permitted motif if and only if each σ_i is a permitted motif. The index formulas also follow directly from Theorems 11 to 13. Note that in the case of the disjoint union, $\widehat{\sigma} = [N]$ and $|\widehat{\sigma}| = N$ because we assume here that each σ_i is nonempty. \square

Observe that Theorems 11 to 13 show that for any disjoint, clique, or cyclic union, $FP(G)$ can only depend on the parameters ε and δ via the $FP(G_i)$. Thus, if $FP(G_i)$ is parameter independent for all i , then $FP(G)$ is also parameter independent.

Corollary 5. *Let G be a disjoint, clique, or cyclic union of components G_1, \dots, G_N with vertex sets τ_1, \dots, τ_N . Then*

$$FP(G) \text{ is parameter independent} \Leftrightarrow FP(G_i) \text{ is parameter independent} \\ \text{for all } i \in [N].$$

In particular, if $|\tau_i| \leq 4$ for all $i \in [N]$, then $FP(G)$ is parameter independent.

Note that the last statement follows from Theorem 6.

5.4 Survival Rules for Disjoint Unions and Clique Unions. Next we consider disjoint, clique, and cyclic unions that are embedded in a larger graph, G , which is not assumed to have any special composite structure.




Figure 13: (A) A cyclic union that does not survive the addition of node k , by graphical domination. (B) A cyclic union that does survive the addition of node k , by uniform in-degree.

Proposition 2 shows that a necessary condition for a disjoint union σ to survive in the larger graph is that every component σ_i is itself a surviving fixed point support. In contrast, Proposition 3 shows that a sufficient condition for a clique union fixed point to survive in G is that at least one component σ_i survives the addition of the other vertices in G . The proofs of these results rely on general domination, and so we save them for section 6.5.

Proposition 2 (*survival of disjoint union*). Let $G|_\sigma$ be a subgraph of G that is a disjoint union of components $G|_{\sigma_1}, \dots, G|_{\sigma_N}$. If $\sigma_i \notin FP(G)$ for any $i \in [N]$, then $\sigma \notin FP(G)$.

Proposition 3 (*survival of clique union*). Let σ be a permitted motif that is a clique union of components $G|_{\sigma_1}, \dots, G|_{\sigma_N}$ inside a larger graph G . If for each $k \in [n] \setminus \sigma$, there exists an $i \in [N]$ such that $\sigma_i \in FP(G|_{\sigma_i \cup k})$, then $\sigma \in FP(G)$.

Propositions 2 and 3 gave simple survival conditions for the composite graph solely in terms of the survival of the individual components. Unfortunately there is no analogous result for cyclic unions because the order of the surviving/non-surviving components within the cyclic union can have an impact on the survival of the composite graph. Specifically consider the graphs in Figure 13, which are cyclic unions with outgoing edges to an external node k . For both of these graphs, the components of the cyclic union are two 2-cliques and two singletons, and in both, the singleton components are permitted motifs that do not survive in the full graph since they have outgoing edges to node k . The only difference between the two graphs is the order in which those components were inserted in the cyclic union. However, the cyclic union does not survive in the full graph in Figure 13A, while it does survive in the full graph in Figure 13B. To see this, note that in panel A, node k outside-in graphically dominates the node to its left, and thus the cyclic union does not survive. In panel B, the cyclic union has uniform in-degree 2 and node k receives only two edges from it; thus, by Rule 2, the cyclic union survives. This contrast shows that we cannot hope for a survival rule for cyclic unions that relies only on knowing the survival of individual components.




Figure 14: Bidirectional simply-added split. In this graph, τ is simply-added to ω and vice versa. Thus, τ is composed of two classes of nodes: projectors onto ω (top dark gray region) and nonprojectors onto ω (bottom light gray region). Similarly, ω can be decomposed into projectors and nonprojectors onto τ . The thick arrows indicate that every node of a given region sends an edge to every node in the other region. The edges within τ and ω can be arbitrary.

5.5 Bidirectional Simply-Added Splits. A key feature of composite graphs that allowed us to prove the previous results is that for each component, G_i , the rest of the graph is simply-added onto it. Furthermore, in the case of disjoint unions and clique unions, every subset of components has the property that the rest of the graph is simply-added onto it, and vice versa. This is also true of all composite graphs with only two components, but is not a general feature of composite graphs. For example, cyclic unions do not have this property: no single component is simply-added onto the rest of the graph.

Here we consider graphs G that have a bidirectional simply-added split, meaning that $[n] = \tau \cup \omega$, where ω is simply-added onto τ and τ is simply-added onto ω . Note that although every composite graph with $N = 2$ components satisfies this property, G need not be such a graph (see Figure 14). For example, τ could send both projectors and nonprojectors to ω , and vice versa, so that $G|_{\tau}$ and $G|_{\omega}$ are not valid components of a composite graph.

The following theorem will use some new notation. Let G be a graph on n nodes. For any $\omega \subseteq [n]$, let S_{ω} denote the fixed point supports of $G|_{\omega}$ that survive to be fixed points of G , and let D_{ω} denote the non-surviving (dying) fixed points:

$$S_{\omega} \stackrel{\text{def}}{=} \text{FP}(G|_{\omega}) \cap \text{FP}(G), \quad \text{and} \quad D_{\omega} \stackrel{\text{def}}{=} \text{FP}(G|_{\omega}) \setminus S_{\omega}.$$

Recall that by Corollary 2, in order to check if $\sigma \in \text{FP}(G|_{\omega})$ survives to $\text{FP}(G)$, we need only to check that $\sigma \in \text{FP}(G|_{\omega \cup k})$ for each $k \notin \omega$. In the case where G has a simply-added split, $[n] = \tau \cup \omega$, with ω simply-added onto τ , each $k \notin \omega$ receives precisely the same edges from ω , and so we need only check that $\sigma \in \text{FP}(G|_{\omega \cup k})$ for a single $k \in \tau$. In this case,

$$S_{\omega} = \{\sigma \subseteq \omega \mid \sigma \in \text{FP}(G|_{\omega \cup k})\},$$

for any choice of $k \in \tau$.

Theorem 14. *Let G be a graph with bidirectional simply-added split $[n] = \tau \cup \omega$. For any nonempty $\sigma \subseteq [n]$, let $\sigma = \sigma_\tau \cup \sigma_\omega$ where $\sigma_\tau = \sigma \cap \tau$ and $\sigma_\omega = \sigma \cap \omega$. Then, $\sigma \in \text{FP}(G)$ if and only if one of the following holds:*

- i. $\sigma_\tau \in S_\tau \cup \emptyset$ and $\sigma_\omega \in S_\omega \cup \emptyset$ or
- ii. $\sigma_\tau \in D_\tau$ and $\sigma_\omega \in D_\omega$.

Moreover, if $\sigma \in \text{FP}(G)$ and both σ_τ and σ_ω are nonempty, then its index is given by

$$\text{idx}(\sigma) = \begin{cases} -\text{idx}(\sigma_\tau)\text{idx}(\sigma_\omega), & \text{if } \sigma_\tau \in S_\tau, \sigma_\omega \in S_\omega \\ \text{idx}(\sigma_\tau)\text{idx}(\sigma_\omega), & \text{if } \sigma_\tau \in D_\tau, \sigma_\omega \in D_\omega \end{cases}.$$

Otherwise, $\sigma = \sigma_\tau$ or σ_ω , and has the same index.

Proof. First we consider the case where σ_τ or σ_ω is empty. Without loss of generality, suppose $\sigma_\tau = \emptyset$. Then $\sigma = \sigma_\omega$, and clearly $\sigma \in \text{FP}(G)$ if and only if $\sigma_\omega \in S_\omega$.

Now suppose $\sigma \subseteq [n]$ with both σ_τ and σ_ω nonempty. Since G has a bidirectional simply-added split along (τ, ω) , using Theorem 3, we see that for any $i \in [n]$, we have

$$s_i^\sigma = \frac{1}{\theta} s_i^{\sigma_\tau} s_i^{\sigma_\omega}.$$

Furthermore, for any $i \in \tau$, all $s_i^{\sigma_\omega}$ have the same value, and so $s_i^\sigma = \alpha s_i^{\sigma_\tau}$. Similarly, for any $j \in \omega$, we have $s_j^\sigma = \beta s_j^{\sigma_\tau}$. Thus, the relative signs of $s_i^{\sigma_\tau}$ across $i \in \tau$ are preserved in the s_i^σ , and the same for the relative signs of $s_j^{\sigma_\omega}$ across $j \in \omega$. Hence, $\sigma \in \text{FP}(G)$ if and only if $\sigma_\tau \in \text{FP}(G|_\tau)$, $\sigma_\omega \in \text{FP}(G|_\omega)$, and $\text{sgn}(s_i^\sigma) = \text{sgn}(s_j^\sigma)$ for any $i \in \sigma_\tau$ and $j \in \sigma_\omega$.

To see when the above signs agree, observe that $\sigma_\tau \in \text{FP}(G|_\tau)$ implies $\text{sgn}(s_i^{\sigma_\tau}) = \text{idx}(\sigma_\tau)$ for all $i \in \sigma_\tau$, and similarly $\text{sgn}(s_j^{\sigma_\omega}) = \text{idx}(\sigma_\omega)$ for all $j \in \sigma_\omega$. Therefore, $\text{sgn}(s_i^\sigma) = \text{sgn}(s_j^\sigma)$ if and only if

$$\text{sgn}(s_i^{\sigma_\omega})\text{idx}(\sigma_\tau) = \text{sgn}(s_j^{\sigma_\tau})\text{idx}(\sigma_\omega). \quad (5.1)$$

However, $\text{sgn}(s_i^{\sigma_\omega})$ and $\text{sgn}(s_j^{\sigma_\tau})$ depend on whether σ_ω and σ_τ , respectively, survive to fixed points of G . To track this, we define $\chi(\sigma_\tau) = 1$ if $\sigma_\tau \in S_\tau$, and $\chi(\sigma_\tau) = -1$ if $\sigma_\tau \in D_\tau$. (Note that $\sigma_\tau \in \text{FP}(G|_\tau)$ implies $\sigma_\tau \in S_\tau \dot{\cup} D_\tau$.) In particular, since $j \notin \sigma_\tau$, we see that $\text{sgn}(s_j^{\sigma_\tau})$ agrees with $\text{idx}(\sigma_\tau)$ if and only if $\sigma_\tau \in D_\tau$. Thus, we can write $\text{sgn}(s_j^{\sigma_\tau}) = -\chi(\sigma_\tau)\text{idx}(\sigma_\tau)$ and $\text{sgn}(s_i^{\sigma_\omega}) = -\chi(\sigma_\omega)\text{idx}(\sigma_\omega)$. Plugging this into equation 5.1, we see that $\text{sgn}(s_i^\sigma) = \text{sgn}(s_j^\sigma)$ if and only if

$$-\chi(\sigma_\omega)\text{idx}(\sigma_\omega)\text{idx}(\sigma_\tau) = -\chi(\sigma_\tau)\text{idx}(\sigma_\tau)\text{idx}(\sigma_\omega),$$

which holds if and only if $\chi(\sigma_\tau) = \chi(\sigma_\omega)$. In other words, we can conclude that $\sigma \in \text{FP}(G)$ if and only if both $\sigma_\tau \in S_\tau$ and $\sigma_\omega \in S_\omega$, or both $\sigma_\tau \in D_\tau$ and $\sigma_\omega \in D_\omega$, as desired. Finally, the index formulas follow from observing that if $\sigma \in \text{FP}(G)$, then $\text{idx}(\sigma) = -\chi(\sigma_\tau)\text{idx}(\sigma_\tau)\text{idx}(\sigma_\omega)$. \square

It is worth noting that Theorem 14 allows us to recover the results on disjoint unions and clique unions as a special case. If τ and ω contain only nonprojectors, then $\tau \cup \omega$ is a disjoint union, where every fixed point of each subset survives: $S_\tau = \text{FP}(G|_\tau)$, $S_\omega = \text{FP}(G|_\omega)$, and $D_\tau = D_\omega = \emptyset$. Thus, Theorem 14 shows that the fixed points of a disjoint union are all the fixed points of the individual components and unions of these. If τ and ω contain only projectors, then $\tau \cup \omega$ is a clique union. In this case, every fixed point of each subset does not survive (dies) because it has a target, and so $D_\tau = \text{FP}(G|_\tau)$, $D_\omega = \text{FP}(G|_\omega)$, and $S_\tau = S_\omega = \emptyset$. Therefore, Theorem 14 shows that the fixed points of a clique union are solely the unions of fixed points from every component.

6 Domination

In this section, we introduce a more general form of domination, which is broader than the concept of graphical domination introduced in section 3.3. Domination applies to all competitive and nondegenerate TLNs with uniform external inputs, so that $b_i = b_j = \theta$ for all $i, j \in [n]$. Furthermore, while graphical domination is insufficient to determine all permitted and forbidden motifs of a CTLN (see appendix section A.3), general domination precisely characterizes all fixed point supports not only for CTLNs but also for TLNs. In particular, our main result on domination, Theorem 15, gives an alternative to Theorem 2 (sign conditions).

As we will see later in this section, this broader form of domination is not practical for explicit computations, but provides a useful technical tool for proving results about fixed points without appealing to the signs of the s_i^σ . In particular, domination will allow us to prove Theorems 4 (graphical domination) and 5 (uniform in-degree), as well as Propositions 2 and 3 (survival rules for disjoint and clique unions).

6.1 General Domination. Let (W, θ) be a TLN with uniform inputs θ , and define $\tilde{W} = -I + W$. The quantities of interest for general domination are the sums:

$$w_j^\sigma \stackrel{\text{def}}{=} \sum_{i \in \sigma} \tilde{W}_{ji} |s_i^\sigma|. \quad (6.1)$$




Figure 15: The butterfly graph for example 6.

In contrast to Theorem 2 (sign conditions), where the signs of the s_i^σ were essential to determining fixed point supports, here we completely discard the signs and use only the absolute values $|s_i^\sigma|$.

The definition of the w_j^σ may seem mysterious at first. Before defining domination or stating our main theorem about it, we will work through an example that shows how the w_j^σ values encode information about the fixed point supports of a TLN. Two main observations will emerge from this example. First, if $\sigma \in \text{FP}(W, \theta)$, then $w_j^\sigma = w_k^\sigma$ for all $j, k \in \sigma$. In other words, the w_j^σ precisely match for all nodes inside a permitted motif. Second, if $\sigma \in \text{FP}(W, \theta)$ and there is some $k \notin \sigma$, then for $j \in \sigma$, we have $w_j^\sigma > w_k^\sigma$. So the values of w_j^σ inside the fixed point support are all equal to each other and greater than the values of w_k^σ for nodes *outside*.

Example 6. Consider a CTLN whose graph G is the butterfly graph in Figure 15, which is studied in detail in appendix section A.3. By appealing to earlier graph rules and parity, it is straightforward to see that $\text{FP}(G) = \{123, 234, 1234\}$. In this example, we explore the values of w_j^σ for different subsets of vertices in this graph.

First, consider $\sigma = \{1, 2, 3, 4\}$, which we already know is a permitted motif. In appendix section A.3, we compute the following values for the butterfly graph:

$$\begin{aligned} s_1^\sigma &= s_4^\sigma = -\delta\theta(\varepsilon^2 + \varepsilon\delta + \delta^2), & s_2^\sigma &= -\delta\theta(2\varepsilon^2 + 2\varepsilon\delta + \delta^2), \\ s_3^\sigma &= -\delta\theta(2\varepsilon^2 + 3\varepsilon\delta + 2\delta^2). \end{aligned}$$

Using these calculations, we can compute the w_j^σ values—for example,

$$\begin{aligned} w_1^\sigma &= \sum_{i \in \sigma} \tilde{W}_{1i} |s_i^\sigma| = (-1)|s_1^\sigma| + (-1 - \delta)|s_2^\sigma| + (-1 + \varepsilon)|s_3^\sigma| + (-1 - \delta)|s_4^\sigma| \\ &= \delta\theta(2\varepsilon^3 - 6\varepsilon^2 - 7\varepsilon\delta - \varepsilon\delta^2 - 5\delta^2 - 2\delta^3). \end{aligned}$$

By symmetry, it is clear that $w_4^\sigma = w_1^\sigma$, but in fact, all the w_j^σ values are equal across $j \in \sigma$. Although it is not obvious that w_2^σ should match w_1^σ , it does:

$$\begin{aligned} w_2^\sigma &= \sum_{i \in \sigma} \tilde{W}_{2i} |s_i^\sigma| = (-1)|s_2^\sigma| + (-1 + \varepsilon)(|s_1^\sigma| + |s_4^\sigma|) + (-1 - \delta)|s_3^\sigma| \\ &= \delta\theta(2\varepsilon^3 - 6\varepsilon^2 - 7\varepsilon\delta - \varepsilon\delta^2 - 5\delta^2 - 2\delta^3). \end{aligned}$$

Theorem 15 will show that this is the hallmark of permitted motifs.

Next consider $\tau = \{1, 2, 3\}$. Since τ is a 3-cycle, the s_i^τ values can be obtained from graph 8 in Figure 18, giving us $s_i^\tau = \theta(\varepsilon^2 + \varepsilon\delta + \delta^2)$ for all $i \in \tau$. Then, for all $j \in \tau$,

$$\begin{aligned} w_j^\tau &= \sum_{i \in \sigma} \tilde{W}_{ji} |s_i^\sigma| = \theta(\varepsilon^2 + \varepsilon\delta + \delta^2)(-1 + (-1 + \varepsilon) + (-1 - \delta)) \\ &= \theta(\varepsilon^2 + \varepsilon\delta + \delta^2)(-3 + \varepsilon - \delta). \end{aligned}$$

In addition,

$$\begin{aligned} w_4^\tau &= \sum_{i \in \sigma} \tilde{W}_{4i} |s_i^\sigma| = \theta(\varepsilon^2 + \varepsilon\delta + \delta^2)((-1 + \varepsilon) + 2(-1 - \delta)) \\ &= \theta(\varepsilon^2 + \varepsilon\delta + \delta^2)(-3 + \varepsilon - 2\delta). \end{aligned}$$

Notice that $w_j^\tau > w_4^\tau$ for all $j \in \tau$. In fact, the inequality that $w_j^\tau > w_k^\tau$ for all $j \in \tau$ and $k \notin \tau$ must be satisfied for any permitted motif τ to survive as a fixed point of a larger graph G . This observation, together with the fact that within a permitted motif all the values w_j^σ match, is captured below in Theorem 15.

In order to compute the values w_j^σ in example 6, we used our precomputed values for s_i^σ . However, once we know the s_i^σ , we can simply apply Theorem 2 (sign conditions) and be done. In particular, it is not practical to explicitly compute the w_j^σ in order to check whether σ is a fixed point support. The true value of introducing the w_j^σ is that arguments can be made about their relative values without knowing the signs of the s_i^σ , and this in turn can be used to determine whether $\sigma \in \text{FP}(W, \theta)$. With this motivation, we now define *domination*, which is a generalization of graphical domination (introduced in section 3.3).

Definition 10 (*domination*). Consider a TLN (W, θ) on n neurons, let $\sigma \subseteq [n]$ be nonempty, and let w_j^σ be defined as in equation 6.1. For any $j, k \in [n]$, we say that

- k dominates j with respect to σ , and write $k >_\sigma j$, whenever $w_k^\sigma > w_j^\sigma$.
- k is equivalent to j with respect to σ , and write $k \sim_\sigma j$, whenever $w_k^\sigma = w_j^\sigma$.




Figure 16: Graph for example 7.

- k is not dominated by j with respect to σ , and write $k \geq_{\sigma} j$, whenever $w_k^{\sigma} \geq w_j^{\sigma}$.

Moreover, if $j \sim_{\sigma} k$ for all $j, k \in \sigma$, then we say that σ is domination-free.

The domination relation satisfies nice properties. Clearly, if $\ell >_{\sigma} k$ and $k >_{\sigma} j$, then $\ell >_{\sigma} j$. Furthermore, if $k >_{\sigma} j$, then we cannot also have $j >_{\sigma} k$. We thus see that $>_{\sigma}$ is transitive and antisymmetric but not reflexive, while \geq_{σ} is transitive, antisymmetric, and reflexive. This makes $>_{\sigma}$ a strict partial order and \geq_{σ} a partial order. Incomparable elements under $>_{\sigma}$ belong to equivalence classes of the equivalence relation \sim_{σ} . It is easy to see that $>_{\sigma}$, \geq_{σ} , and \sim_{σ} interact just as the usual “ $>$ ”, “ \geq ” and “ $=$ ” do. In particular, if $\ell >_{\sigma} k$ and $k \sim_{\sigma} j$, then $\ell >_{\sigma} j$.

We can now give our second characterization of $\text{FP}(W, \theta)$, for the case of uniform inputs $b_i = \theta > 0$ for each $i \in [n]$. Recall from definition 3 that σ is a permitted motif if $\sigma \in \text{FP}(W_{\sigma}, \theta)$ and σ is a forbidden motif otherwise. In order to have $\sigma \in \text{FP}(W, \theta)$, we must have that σ is a permitted motif that survives as a fixed point support in the full network.

Theorem 15 (general domination). *Let (W, θ) be a TLN, and let $\sigma \subseteq [n]$. Then*

$$\sigma \text{ is a permitted motif} \Leftrightarrow \sigma \text{ is domination free.}$$

If σ is a permitted motif, then $\sigma \in \text{FP}(W, \theta)$ if and only if for each $k \notin \sigma$, there exists $j \in \sigma$ such that $j >_{\sigma} k$.

In addition to telling us that permitted motifs are domination free, Theorem 15 states that permitted motifs survive precisely when every node outside the motif is (inside-out) dominated by some node inside. This is precisely what we saw in example 6. Namely, the w_j^{τ} and w_j^{σ} all matched for $j \in \tau$ or $j \in \sigma$, respectively. Moreover, for $j \in \tau$ and $k \notin \tau$, we saw that $w_j^{\tau} > w_k^{\tau}$, and thus $j >_{\tau} k$, consistent with the fact that $\tau \in \text{FP}(G)$.

Before presenting the proof of Theorem 15, in the next section, we illustrate it with another example. Let G be the graph in Figure 16; consider $\tau = \{1, 2, 3, 4\}$. Since τ is uniform in-degree (with $d = 1$) and node 5 receives more than d edges from τ , Theorem 5 guarantees that τ is a permitted

motif that does not survive in the full graph, and so we know that $\tau \in \text{FP}(G|_\tau)$ but $\tau \notin \text{FP}(G)$. However, this result cannot be obtained from graphical domination. In fact, the proof of Theorem 5 relies on general domination. In the following example, we will show that $\tau \notin \text{FP}(G)$ directly, using Theorem 15 (domination). This previews the proof of Theorem 5 in section 6.4.

Example 7. Let G be the graph in Figure 16, and let $\tau = \{1, 2, 3, 4\}$. Observe that for each $i \in \tau$, we can compute $s_i^\tau = \det((I - W_\tau)_i; \theta \mathbf{1}) = \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2)$. Then for all $j \in \tau$, we have

$$\begin{aligned} w_j^\tau &= \sum_{i \in \tau} \tilde{W}_{ji} |s_i^\tau| = \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2) \sum_{i \in \tau} \tilde{W}_{ji} \\ &= \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2)(-4 + \varepsilon - 2\delta). \end{aligned}$$

Since the w_j^τ have the same value for all $j \in \tau$, we see that τ is domination-free, and hence is permitted. Next,

$$\begin{aligned} w_5^\tau &= \sum_{i \in \tau} \tilde{W}_{5i} |s_i^\tau| = \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2) \sum_{i \in \tau} \tilde{W}_{5i} \\ &= \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2)(-4 + 2\varepsilon - 2\delta) > w_j^\tau, \end{aligned}$$

for $j \in \tau$. Thus, $5 \succ_\tau j$, and so $\tau \notin \text{FP}(G)$ by Theorem 15.

Now consider $\sigma = \{1, 2, 3, 4, 5\}$. Observe that node 5 is a nonprojector onto τ , and thus by Corollary 4, $s_i^\sigma = -\delta s_i^\tau = -\varepsilon\delta\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2)$ for all $i \in \tau$. Computing $s_5^\sigma = \det((I - W_\sigma)_5; \theta \mathbf{1})$, we obtain $s_5^\sigma = \varepsilon^2\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2)$. Then for all $j \in \sigma \setminus 5$, we have

$$w_j^\sigma = \sum_{i \in \sigma \setminus 5} \tilde{W}_{ji} |s_i^\sigma| + (-1 - \delta) |s_5^\sigma| = \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2)(-\varepsilon - 4\delta - 2\delta^2),$$

while

$$w_5^\sigma = \sum_{i \in \sigma \setminus 5} \tilde{W}_{5i} |s_i^\sigma| + (-1) |s_5^\sigma| = \varepsilon\theta(\varepsilon^2 + 2\varepsilon\delta + 2\delta^2)(-\varepsilon + 2\varepsilon\delta - 4\delta - 2\delta^2).$$

Since $w_5^\sigma \neq w_j^\sigma$ for $j, 5 \in \sigma$, we see that σ is not domination-free, and hence σ is a forbidden motif.⁶

6.2 Proof of Theorem 15 (General Domination). Here we assume $b_i = \theta$ for all $i \in [n]$. For constant $\theta > 0$, the collection of fixed point supports

⁶Note that we also could have concluded that σ was forbidden by Rule 8 since it is the union of sink with a non-surviving fixed point.

is independent of the value of θ , and so we drop θ from the notation and denote the set of fixed point supports simply as $\text{FP}(W)$.

Note that $s_i^\sigma = \det((I - W_\sigma)_j; \theta \mathbf{1})$ for each $i \in \sigma$, and equation 2.6 implies

$$-s_k^\sigma + \sum_{i \in \sigma \setminus \{k\}} W_{ki} s_i^\sigma = -s_j^\sigma + \sum_{i \in \sigma \setminus \{j\}} W_{ji} s_i^\sigma \quad \text{for all } j, k \in [n],$$

where we have used the fact that $b_k = b_j$ and $W_{kk} = W_{jj} = 0$. If we denote $\tilde{W} = -I + W$, then the terms s_j^σ, s_k^σ can be absorbed into the sum when $j, k \in \sigma$. This allows us to write

$$(\chi_\sigma(k) - 1)s_k^\sigma + \sum_{i \in \sigma} \tilde{W}_{ki} s_i^\sigma = (\chi_\sigma(j) - 1)s_j^\sigma + \sum_{i \in \sigma} \tilde{W}_{ji} s_i^\sigma \quad \text{for all } j, k \in [n], \quad (6.2)$$

where $\chi_\sigma(k) = 1$ if $k \in \sigma$ and $\chi_\sigma(k) = 0$ if $k \notin \sigma$. In particular, if $j, k \in \sigma$, the above expression reduces to

$$\sum_{i \in \sigma} \tilde{W}_{ki} s_i^\sigma = \sum_{i \in \sigma} \tilde{W}_{ji} s_i^\sigma, \quad \text{for } j, k \in \sigma. \quad (6.3)$$

Equations 6.2 and 6.3 are true in general, irrespective of whether σ is a fixed point support.

Now assume σ is a permitted motif, so that $\sigma \in \text{FP}(W_\sigma)$. In this case, all the s_i^σ for $i \in \sigma$ must have the same sign, and so equations 6.2 and 6.3 continue to hold if each s_i^σ is replaced by $|s_i^\sigma|$. In fact, in the case of equation 6.3, the converse is also true: if equality holds after replacing s_i^σ with $|s_i^\sigma|$, then $\sigma \in \text{FP}(W_\sigma)$. Using the notation $w_j^\sigma = \sum_{i \in \sigma} \tilde{W}_{ji} |s_i^\sigma|$, as in equation 6.1, we have the following lemma.

Lemma 13. $\sigma \in \text{FP}(W_\sigma) \Leftrightarrow w_j^\sigma = w_k^\sigma$ for all $j, k \in \sigma$. (I.e., $\sigma \in \text{FP}(W_\sigma) \Leftrightarrow j \sim_\sigma k$ for all $j, k \in \sigma$.)

Proof (\Rightarrow) This direction was already established in the arguments above. (\Leftarrow) To see the converse, suppose $w_j^\sigma = w_k^\sigma$ for all $j, k \in \sigma$. Let $\alpha = -w_j^\sigma = -\sum_{i \in \sigma} \tilde{W}_{ji} |s_i^\sigma| = \sum_{i \in \sigma} (I - W)_{ji} |s_i^\sigma|$. Note that $\alpha > 0$, and let $v = (|s_i^\sigma|)_{i \in \sigma}$ be the column vector whose entries are $|s_i^\sigma|$ for each $i \in \sigma$. It follows that $(I - W_\sigma)v = \alpha \mathbf{1}_\sigma$, and so for $x_\sigma = \frac{\alpha}{\alpha} v$, we have $(I - W_\sigma)x_\sigma = \theta \mathbf{1}_\sigma$. Since x_σ has strictly positive entries, it is a fixed point of the network (W_σ, θ) with support σ . Thus, $\sigma \in \text{FP}(W_\sigma)$. \square

Lemma 13 can be restated as saying that $\sigma \in \text{FP}(W_\sigma)$ if and only if σ is domination-free. Interestingly, the nondegeneracy condition on TLNs guarantees that if σ is a permitted motif, then a node outside σ can never have the same w_j^σ value as one inside σ , and so one must dominate the other.

Lemma 14. *Let $\sigma \in FP(W_\sigma)$. If $j \in \sigma$ and $k \notin \sigma$, then either $k >_\sigma j$ or $j >_\sigma k$ (i.e., we cannot have $j \sim_\sigma k$).*

Proof. It follows from equation 6.2 that if $j \in \sigma$ and $k \notin \sigma$, then

$$s_k^\sigma = \sum_{i \in \sigma} \tilde{W}_{ki} s_i^\sigma - \sum_{i \in \sigma} \tilde{W}_{ji} s_i^\sigma,$$

and so $\text{sgn } s_k^\sigma = \text{sgn}(\sum_{i \in \sigma} \tilde{W}_{ki} s_i^\sigma - \sum_{i \in \sigma} \tilde{W}_{ji} s_i^\sigma)$. If $\sigma \in FP(W_\sigma)$, then all s_i^σ for $i \in \sigma$ have the same sign. Replacing s_i^σ with $|s_i^\sigma|$ yields

$$\text{sgn } s_k^\sigma = \pm \text{sgn}(w_k^\sigma - w_j^\sigma).$$

By the assumption of nondegeneracy of the TLN, we know that $\text{sgn } s_k^\sigma \neq 0$, and so either $w_k^\sigma > w_j^\sigma$ or $w_j^\sigma > w_k^\sigma$. \square

The next lemma tells us when fixed points survive the addition of a single node.

Lemma 15. *Suppose $\sigma \in FP(W_\sigma)$, and $k \notin \sigma$. Then $\sigma \in FP(W_{\sigma \cup \{k\}})$ if and only if $j >_\sigma k$ for some $j \in \sigma$. If $\sigma \notin FP(W_{\sigma \cup \{k\}})$, then $k >_\sigma j$ for all $j \in \sigma$.*

Proof. Let $\sigma \in FP(W_\sigma)$, $k \notin \sigma$, and $j \in \sigma$. Recall from the proof of lemma 14 that $\text{sgn } s_k^\sigma = \text{sgn}(\sum_{i \in \sigma} \tilde{W}_{ki} s_i^\sigma - \sum_{i \in \sigma} \tilde{W}_{ji} s_i^\sigma)$. By Theorem 2 and Corollary 2, $\sigma \in FP(W_{\sigma \cup \{k\}})$ if and only if $\text{sgn } s_k^\sigma = -\text{sgn } s_j^\sigma$. If $\text{sgn } s_i^\sigma = +1$ for each $i \in \sigma$, then replacing s_i^σ with $|s_i^\sigma|$ in the sums reveals that $\sigma \in FP(W_{\sigma \cup \{k\}})$ if and only if $j >_\sigma k$. Similarly, if $\text{sgn } s_i^\sigma = -1$, we also have that $\sigma \in FP(W_{\sigma \cup \{k\}})$ if and only if $j >_\sigma k$. On the other hand, if $\sigma \notin FP(W_{\sigma \cup \{k\}})$, then we must have $k >_\sigma j$ by Lemma 14. \square

Combining these results with Corollary 2, we obtain the proof of Theorem 15.

Proof of Theorem 15 (domination). The first statement, that $\sigma \in FP(W_\sigma)$ if and only if σ is domination-free, follows directly from Lemma 13. Next, recall that by corollary 2, we have $\sigma \in FP(W)$ if and only if $\sigma \in FP(W_\sigma)$ and $\sigma \in FP(W_{\sigma \cup \{k\}})$ for all $k \notin \sigma$. Applying Lemma 15, we can conclude that for any permitted motif σ , we have $\sigma \in FP(W)$ if and only if for each $k \notin \sigma$, there exists a $j \in \sigma$ such that $j >_\sigma k$. \square

We end this section with a lemma that collects some key facts relating domination to fixed point supports and parallels Theorem 4 (graphical domination). In fact, we will use this in the next section to prove Theorem 4.

Lemma 16 (domination). *Let (W, θ) be a TLN. Suppose $k >_\sigma j$ for some $j, k \in [n]$:*

- If $j, k \in \sigma$, then $\sigma \notin FP(W_\sigma)$, and thus $\sigma \notin FP(W)$.*
- If $j \in \sigma, k \notin \sigma$, then $\sigma \notin FP(W_{\sigma \cup \{k\}})$, and thus $\sigma \notin FP(W)$.*
- If $j \notin \sigma, k \in \sigma$, and $\sigma \in FP(W_\sigma)$, then $\sigma \in FP(W_{\sigma \cup \{j\}})$.*

Proof. Part a is a direct corollary of Theorem 15, while parts b and c are direct consequences of Lemma 15. \square

6.3 Proof of Theorem 4 (Graphical Domination). Using Lemma 16, it is now straightforward to prove Theorem 4. But first we need to show that graphical domination (as defined in section 3.3) is indeed a special case of the more general domination.

Lemma 17. *Consider a CTLN on n nodes, and let $j, k \in [n]$ and $\sigma \subseteq [n]$. Suppose k graphically dominates j with respect to σ . Then $k >_\sigma j$.*

Proof. Recall that if k graphically dominates j with respect to σ , then $\sigma \cap \{j, k\} \neq \emptyset$. Moreover, three conditions hold: (1) for each $i \in \sigma \setminus \{j, k\}$, if $i \rightarrow j$, then $i \rightarrow k$; (2) if $j \in \sigma$, then $j \rightarrow k$; and (3) if $k \in \sigma$, then $k \not\rightarrow j$.

Next, recall that in a CTLN, we have $\tilde{W}_{ji} = -1 + \varepsilon$ if $i \rightarrow j$, $\tilde{W}_{ji} = -1 - \delta$ if $i \not\rightarrow j$, and $\tilde{W}_{ii} = -1$. Condition 1 thus implies that $\tilde{W}_{ji} \leq \tilde{W}_{ki}$ for each $i \in \sigma \setminus \{j, k\}$. If $j \in \sigma$, condition 2 gives $\tilde{W}_{jj} < \tilde{W}_{kj}$, while if $k \in \sigma$, condition 3 implies $\tilde{W}_{jk} < \tilde{W}_{kk}$. Putting these together, we see that

$$w_j^\sigma = \sum_{i \in \sigma} \tilde{W}_{ji} |s_i^\sigma| < \sum_{i \in \sigma} \tilde{W}_{ki} |s_i^\sigma| = w_k^\sigma,$$

where the inequality is strict because at least one of j or k is in σ , and $\tilde{W}_{ji} < \tilde{W}_{ki}$ for $i = j$ or $i = k$. Since $w_k^\sigma > w_j^\sigma$, it follows from the definition that $k >_\sigma j$. \square

We can now prove Theorem 4, which tells us how to use graphical domination in order to rule in and rule out various fixed point supports.

Proof of Theorem 4. First, observe using Lemma 17 that $k >_\sigma j$. Now the statements a, b, and c all follow immediately from parts a, b, and c of Lemma 16. \square

Note that the converse of Lemma 17 is not true: there could still be a domination relationship even if there is no graphical domination. For example, appendix section A.3 shows six graphs in Figure 20 that are forbidden, and thus are not domination-free by Theorem 15. So there must be general domination in each of these graphs, despite the absence of graphical domination.

Note also that unlike graphical domination, general domination may be parameter dependent, even within the legal range. For example, all the graphs in Figure 19 appendix section A.2 have parameter-dependent domination relationships, and this is reflected by the fact that $\text{FP}(G)$ depends on ε and δ .

6.4 Proof of Theorem 5 (Uniform In-Degree). In this section, we use general domination to prove Theorem 5, giving conditions for when a uniform in-degree subset supports a fixed point. We begin by showing that

a uniform in-degree subset always supports a fixed point in its restricted subgraph and that fixed point has uniform firing rate values.

Lemma 18. *If σ has uniform in-degree d , then $\sigma \in FP(G|_\sigma)$ and the corresponding fixed point x^* is uniform, with values*

$$x_i^* = \frac{\theta}{|\sigma| + \delta(|\sigma| - d - 1) - \varepsilon d} \text{ for each } i \in \sigma.$$

Proof. Because $G|_\sigma$ has uniform in-degree, the row sums of $I - W_\sigma$ are all equal. This implies that the all-ones vector $\mathbf{1}_\sigma$ is an eigenvector of $I - W_\sigma$, with eigenvalue R equal to the row sum. Now consider the vector x^* satisfying $x_i^* = \theta/R$ for each $i \in \sigma$, and $x_k^* = 0$ for each $k \notin \sigma$. Clearly, $(I - W_\sigma)x_\sigma^* = \theta\mathbf{1}_\sigma$, and so x_σ^* is a fixed point of the network restricted to σ . Moreover, since all vertices in $G|_\sigma$ have in-degree d , then each row of W_σ has d terms with value $-1 + \varepsilon$ and $|\sigma| - d - 1$ terms with value $-1 - \delta$. This allows us to compute the row sum as

$$R = 1 - \sum_{i \in \sigma} W_{1i} = |\sigma| + \delta(|\sigma| - d - 1) - \varepsilon d,$$

yielding $x_i^* = \theta/(|\sigma| + \delta(|\sigma| - d - 1) - \varepsilon d)$ for $i \in \sigma$, as desired. Note that $x_i^* > 0$ for the full range of d values, so this fixed point always satisfies fixed point condition 1. \square

Next, we give the survival rule of a uniform in-degree fixed point in terms of domination.

Lemma 19 (*uniform in-degree domination*). *Suppose σ has uniform in-degree d , and suppose $j \in \sigma$ and $k \notin \sigma$. Let $d_k = |\{i \in \sigma \mid i \rightarrow k\}|$ be the number of edges k receives from σ . Then $\sigma \in FP(G|_\sigma)$ and*

- i. $k >_\sigma j$ if $d_k > d$.
- ii. $j >_\sigma k$, if $d_k \leq d$.

Proof. By Cramer’s rule, $s_i^\sigma = \det(I - W_\sigma)x_i^*$ (see Lemma 1 and equation 2.3). Thus, by Lemma 18, $s_i^\sigma = s_j^\sigma$ for all $i, j \in \sigma$, when σ is uniform in-degree. This implies $\sigma \in FP(G|_\sigma)$ and also allows us to factor $|s_i^\sigma|$ out of the sums for checking domination, so that $k >_\sigma j$ if and only if $\sum_{i \in \sigma} \tilde{W}_{ki} > \sum_{i \in \sigma} \tilde{W}_{ji}$. Now observe that $\sum_{i \in \sigma} \tilde{W}_{ki} = d_k(-1 + \varepsilon) + (|\sigma| - d_k)(-1 - \delta)$, while $\sum_{i \in \sigma} \tilde{W}_{ji} = d(-1 + \varepsilon) + (|\sigma| - d - 1)(-1 - \delta) - 1 = d(-1 + \varepsilon) + (|\sigma| - d)(-1 - \delta) + \delta$, since $j \in \sigma$ and $\tilde{W}_{jj} = -1$. In particular, if $d_k = d$, then we have $\sum_{i \in \sigma} \tilde{W}_{ki} < \sum_{i \in \sigma} \tilde{W}_{ji}$, so that $j >_\sigma k$. It is now easy to check that if $d_k > d$, then $k >_\sigma j$, while if $d_k \leq d$, then $j >_\sigma k$. \square

Finally, we combine these results to prove Theorem 5 and prove the stability conditions.

Proof of Theorem 5. By Lemma 18, if σ has uniform in-degree, then it supports a fixed point in $G|_{\sigma}$, and thus $i \sim_{\sigma} j$ for all $i, j \in \sigma$. By Lemma 19, for each $k \notin \sigma$, we have $j >_{\sigma} k$ precisely when $d_k \leq d$. Thus, by Theorem 15 (domination), σ supports a fixed point in $G|_{\sigma \cup k}$ if and only if $d_k \leq d$.

For the stability conditions, recall that the fixed point is stable precisely when all the eigenvalues of $-I + W_{\sigma}$ have negative real part or, equivalently, all the eigenvalues of $I - W_{\sigma}$ have positive real part. First consider $d < |\sigma|/2$. Observe that the uniform in-degree implies that the all-ones vector $\mathbf{1}$ is an eigenvector of $I - W_{\sigma}$, with eigenvalue λ equal to the row sum:

$$\lambda = |\sigma| + (|\sigma| - d - 1)\delta - d\varepsilon = |\sigma| + (|\sigma| - 1)\delta - d(\delta + \varepsilon).$$

When $d < |\sigma|/2$, we have $d \leq (|\sigma| - 1)/2$, and thus $\lambda > |\sigma|$ whenever $|\sigma| > 1$ because $\varepsilon < \delta$. Since the sum of the eigenvalues equals the trace, $\text{Tr}(I - W_{\sigma}) = |\sigma|$, we see that $I - W_{\sigma}$ must have a negative eigenvalue. This implies that the fixed point is unstable.

Next consider when $d = |\sigma| - 1$, so that σ is a clique. In this case, $I - W_{\sigma} = (1 - \varepsilon)\mathbf{1}\mathbf{1}^T + \varepsilon I_{\sigma}$, and so the eigenvalues are $|\sigma|(1 - \varepsilon) + \varepsilon$ and ε . Clearly, these are positive for $0 < \varepsilon < \frac{\delta}{\delta+1} < 1$, so we can conclude that the fixed point is stable. \square

6.5 Domination and Simply-Added Splits. When $\sigma = \tau \cup \omega$, where ω is simply-added to τ , we find that domination relationships with respect to τ are preserved with respect to σ .

Lemma 20. *Let $\sigma = \tau \cup \omega$, where ω is simply-added to τ . Then for any $j, k \in \tau$,*

- i. $k >_{\sigma} j \Leftrightarrow k >_{\tau} j$
- ii. $k \sim_{\sigma} j \Leftrightarrow k \sim_{\tau} j$.

Furthermore, for any $j \in \tau$ and $\ell \notin \sigma$:

iii. *If for all $i \in \omega$ such that $i \rightarrow j$, we also have $i \rightarrow \ell$; then*

$$\ell >_{\tau} j \Rightarrow \ell >_{\sigma} j, \text{ and } \ell \geq_{\tau} j \Rightarrow \ell \geq_{\sigma} j.$$

iv. *If for all $i \in \omega$ such that $i \rightarrow \ell$ we also have $i \rightarrow j$, then*

$$j >_{\tau} \ell \Rightarrow j >_{\sigma} \ell, \text{ and } j \geq_{\tau} \ell \Rightarrow j \geq_{\sigma} \ell.$$

In particular, the lemma tells us that if $\tau \notin \text{FP}(G|_{\tau})$ then $\sigma \notin \text{FP}(G)$.

Proof. Using Theorem 3, we see that for each $i \in \tau$, $|s_i^{\sigma}| = |\alpha||s_i^{\tau}|$, where $\alpha = s_i^{\omega}$, which is constant across $i \in \tau$. In particular, for any fixed $j \in \tau$, we have

$$\begin{aligned} w_j^{\sigma} &= \sum_{i \in \sigma} \tilde{W}_{ji}|s_i^{\sigma}| = \sum_{i \in \tau} \tilde{W}_{ji}|s_i^{\sigma}| + \sum_{i \in \omega} \tilde{W}_{ji}|s_i^{\sigma}| \\ &= \sum_{i \in \tau} \tilde{W}_{ji}|\alpha||s_i^{\tau}| + \sum_{i \in \omega} \tilde{W}_{ji}|s_i^{\sigma}| = |\alpha|w_j^{\tau} + \sum_{i \in \omega} \tilde{W}_{ji}|s_i^{\sigma}|, \end{aligned}$$

where the last sum is identical for all $j \in \tau$ because ω is simply-added to τ . This means that for any $j, k \in \tau$,

$$w_k^\sigma - w_j^\sigma = |\alpha|(w_k^\tau - w_j^\tau).$$

Thus, parts i and ii both hold.

By the same logic, for any $j \in \tau$ and $\ell \notin \sigma$, we have $w_\ell^\sigma = |\alpha|w_\ell^\tau + \sum_{i \in \omega} \tilde{W}_{\ell i} |s_i^\sigma|$ and thus,

$$w_\ell^\sigma - w_j^\sigma = |\alpha|(w_\ell^\tau - w_j^\tau) + \sum_{i \in \omega} (\tilde{W}_{\ell i} - \tilde{W}_{ji}) |s_i^\sigma|.$$

In part iii, we have $\tilde{W}_{\ell i} \geq \tilde{W}_{ji}$ for all $i \in \omega$. Thus, if $\ell >_\tau j$, so that $w_\ell^\tau > w_j^\tau$, then $w_\ell^\sigma - w_j^\sigma > 0$ and hence $\ell >_\sigma j$. Similarly, we see that $\ell \geq_\tau j \Rightarrow \ell \geq_\sigma j$. Finally, in part iv, we have $\tilde{W}_{\ell i} \leq \tilde{W}_{ji}$ for all $i \in \omega$. Thus, if $j >_\tau \ell$, so that $w_j^\tau > w_\ell^\tau$, then $w_\ell^\sigma - w_j^\sigma < 0$ and thus $j >_\sigma \ell$. Similarly, $j \geq_\tau \ell \Rightarrow j \geq_\sigma \ell$. \square

Lemma 20 shows that identifying a simply-added split of σ can be useful for identifying domination relationships with respect to τ that lift to domination relationships with respect to σ . This is particularly useful in composite graphs, since for any component of a composite graph, the rest of the graph is simply-added to that component. In particular, we use Lemma 20 to prove the partial survival rules for disjoint unions and clique unions embedded in a larger graph.

Proof of Proposition 2 (survival of disjoint union). Suppose there exists $i \in [N]$ such that $\sigma_i \notin \text{FP}(G)$. If $\sigma_i \notin \text{FP}(G|_{\sigma_i})$, then by Lemma 7, $\sigma \notin \text{FP}(G)$. Assuming $\sigma_i \in \text{FP}(G|_{\sigma_i})$, the fact that $\sigma_i \notin \text{FP}(G)$ implies there exists $j \in \sigma_i$ and $\ell \notin \sigma_i$ such that $\ell \geq_{\sigma_i} j$ (by Theorem 15). Note, however, that for all $k \in \sigma \setminus \sigma_i$, we have $j >_{\sigma_i} k$ because (by the disjoint union) there are no edges from σ_i to k . Hence $\ell \notin \sigma$, and it suffices by Theorem 15 to show that $\ell \geq_\sigma j$ in order to conclude that $\sigma \notin \text{FP}(G)$. To do this, we will use Lemma 20 for $\sigma = \tau \cup \omega$, where $\tau = \sigma_i$ and $\omega = \sigma \setminus \sigma_i$. Since there are no edges from ω to σ_i , the condition in part iii of Lemma 20 is trivially satisfied, and thus $\ell \geq_{\sigma_i} j$ implies $\ell \geq_\sigma j$, as desired. \square

Proof of Proposition 3 (survival of clique union). Since σ is a permitted motif, for all $j, \ell \in \sigma$ we have $j \sim_\sigma \ell$, by Theorem 15. By Theorem 15, to show $\sigma \in \text{FP}(G)$, we must show that for each $k \in [n] \setminus \sigma$, there exists a $j \in \sigma$ such that $j >_\sigma k$. Fix $k \in [n] \setminus \sigma$. By hypothesis, there exists a σ_i such that $\sigma_i \in \text{FP}(G|_{\sigma_i \cup k})$. Thus, for any $j \in \sigma_i$, we have $j >_{\sigma_i} k$. We will now use Lemma 20 to show that this domination relationship lifts to σ . Following the notation of the lemma, let $\sigma = \tau \cup \omega$ for $\tau = \sigma_i$ and $\omega = \sigma \setminus \sigma_i$, and observe that because σ is a clique union, ω is simply-added to τ . Moreover,

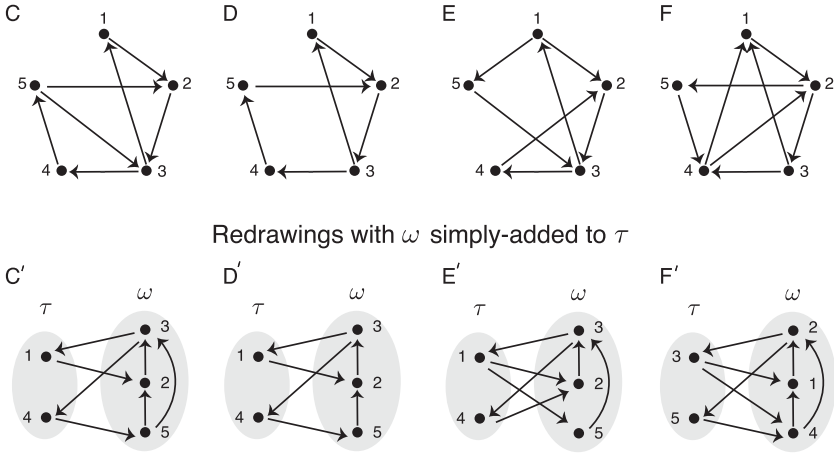


Figure 17: Inferring domination from a simply-added split. (Top) Oriented graphs C–F from Figure 20, which are forbidden without graphical domination. (Bottom) Redrawings of the graphs in terms of components τ and ω , where ω is simply-added to τ .

since $j \in \sigma_i = \tau$ receives edges from all nodes in ω , the condition in part iv of Lemma 20 is trivially satisfied. Thus, $j >_{\sigma_i} k$ implies $j >_{\sigma} k$, as desired. We conclude that $\sigma \in \text{FP}(G|_{\sigma \cup k})$ for each $k \in [n] \setminus \sigma$, and thus $\sigma \in \text{FP}(G)$. \square

It turns out that even when τ is domination-free, it can still be useful to have a simply-added split in order to identify domination relationships in σ . Example 8 illustrates how one can use a simply-added split where τ has uniform in-degree to infer domination relationships in σ from graphical domination that occurs in some “equivalent” graph.

Example 8. Consider the graphs in Figure 17. In appendix section A.3, these graphs are all shown to be forbidden motifs by parity arguments (these are graphs C–F in Figure 20). Here we explore how and when simply-added splits can be used to infer domination relationships to directly show that a graph is forbidden.

Panel C’ in Figure 17 shows a redrawing of graph C in terms of a simply-added split, where $\tau = \{1, 4\}$, $\omega = \{2, 3, 5\}$, and ω is simply-added to τ . Since τ is an independent set, it has uniform in-degree, and thus $s_1^\sigma = s_4^\sigma$. Furthermore, since $\sigma = \{1, 2, 3, 4, 5\}$ is $\tau \cup \omega$, we have that $s_1^\sigma = s_4^\sigma$ by Theorem 3 (simply-added). Thus, any inputs from node 1 to the rest of the graph are equivalent to inputs from node 4 in terms of the contribution of the s_i^σ values (the relevant quantity for computing w_j^σ and identifying domination). Hence, in terms of domination relationships with respect to σ , graph C is equivalent to a graph where the $1 \rightarrow 2$ edge is replaced with a $4 \rightarrow 2$ edge.

In this equivalent graph, we see that node 2 inside-in graphically dominates 5. Thus, in graph C, we have $2 >_{\sigma} 5$, and so σ (i.e., the full graph C) is a forbidden motif since it is not domination-free.

Since graph D differs from C only by the $5 \rightarrow 3$ edge, an identical argument shows that graph D is equivalent (in terms of s_i^{σ} values) to a graph where node 2 graphically dominates 5, and so in graph D, we have $2 >_{\sigma} 5$. Thus, graph D is a forbidden motif.

Observe that graphs E and F can also be decomposed as ω simply-added to τ , where τ has uniform in-degree, and thus the values of s_i^{σ} for $i \in \tau$ are all equal (see Figures 17E' and 17F'). However, there is no equivalent graph with inside-in graphical domination for either of these graphs. To see this in graph E', note that for each pair of nodes $j, k \in \omega$ with $j \rightarrow k$ and $k \not\rightarrow j$ (conditions 2 and 3 of the definition of graphical domination), we have that node j receives more inputs from τ than node k does; thus, condition 1 can never be satisfied in any equivalent graph, and so $k \not>_{\sigma} j$ for any such pair. The same argument can be made for graph F' to show that graphical domination cannot be used to show the motif is forbidden. Thus, to show that these graphs are forbidden motifs, one must explicitly compute the s_i^{σ} or rely on parity arguments as in appendix section A.3.

7 Conclusion

In this work, we have introduced two new characterizations for the fixed points of competitive TLNs: first in terms of sign conditions (see Theorem 2) and later in terms of domination (see Theorem 15). Specializing to CTLNs, we used these tools to prove key theorems on graphical domination (see Theorem 4) and simply-added splits (see Theorem 3), as well as to derive survival rules for uniform in-degree motifs (see Theorem 5). These methods then enabled us to prove a series of graph rules in section 4 that allow one to determine elements of $\text{FP}(G)$ by direct analysis of the graph G . Finally, in section 5, we showed how this graphical calculus can be extended to larger networks of simpler building blocks.

Any conclusions about a network derived from graph rules are automatically parameter independent. Since some CTLNs do have parameter-dependent permitted motifs (see appendix section A.2), we know that graph rules cannot fully determine $\text{FP}(G)$ in all cases. Nevertheless, it is likely that there are many more graph rules we have yet to discover, covering additional cases where motifs are permitted or forbidden in a parameter-independent manner. In particular, the style of graphical analysis illustrated in example 8 seems to hint at a missing graph rule. Moreover, if Conjecture 1 is true, then there are additional building block graph rules of the form given in Theorem 10 that apply to composite graphs with more complicated skeletons.

To what extent do graph rules extend to more general TLNs? We can immediately say, based on the determinant form of the sign conditions in

Theorem 2, that for any CTLN, there must be an open neighborhood in the (W, b) parameter space in which $\text{FP}(W, b) = \text{FP}(G, \varepsilon, \delta)$ —that is, the fixed point supports of all TLNs in this neighborhood match those of the CTLN. This follows from the fact that the s_i^g are all polynomials in the entries of W and b , and thus vary continuously as a function of these parameters. In particular, for purposes of determining fixed points, the constraints on (W, b) imposed by CTLNs are not fine-tuned, and so the inferences we make from graph rules are robust to at least small perturbations of the network parameters. That said, we currently have no theoretical handle on how big these neighborhoods are beyond the fact that they are open sets (and thus full-dimensional).

There are also non-perturbative approaches to generalizing graph rules. In appendix section A.4, we show how to associate a graph to any TLN (W, θ) having uniform input. For $n = 1, 2$, this graph fully characterizes the permitted motifs, just as it does for CTLNs. However, various CTLN graph rules break down at $n = 3$, as illustrated in example 9. Currently, by applying the theory of oriented matroids to hyperplane arrangements associated with TLNs, we are developing weaker versions of the graph rules that do extend to these more general settings.

Finally, a comment on the nonlinearity. The proofs in this letter all rely on the precise form of the threshold-nonlinearity in equation 1.1, as well as on the assumption that W is competitive. Based on our own computational observations, however, we expect that many qualitative aspects of these results should continue to hold for other nonlinearities, and/or less strict assumptions on W (and b). But these questions are beyond the scope of this letter, and so we leave them for future work.

Appendix

A.1 Permitted and Forbidden Motifs of Size $n \leq 3$. Figure 18 shows all directed graphs of size $n \leq 3$. The permitted motifs are labeled with a “p”; all other graphs are forbidden motifs.

A.2 Parameter-Dependent $\text{FP}(G)$ for $n = 5$. In section 4.3, we saw that $\text{FP}(G)$ is independent of parameters ε and δ when $n \leq 4$; however, there are three permitted motifs of size 4 whose survival in a larger graph is parameter dependent. These permitted motifs are reprinted here in Figures 19A to 19C. The rest of the figure shows example graphs of size $n = 5$ whose $\text{FP}(G) = \text{FP}(G, \varepsilon, \delta)$ depends on the choice of ε and δ . This parameter dependence is a result of these graphs containing one or more permitted motifs whose survival is parameter dependent. Note that none of these $n = 5$ graphs is oriented and therefore is not covered by Theorem 7.

A.3 Using Parity to Detect Forbidden Motifs without Graphical Domination. Although inside-in graphical domination tells us that a motif in a

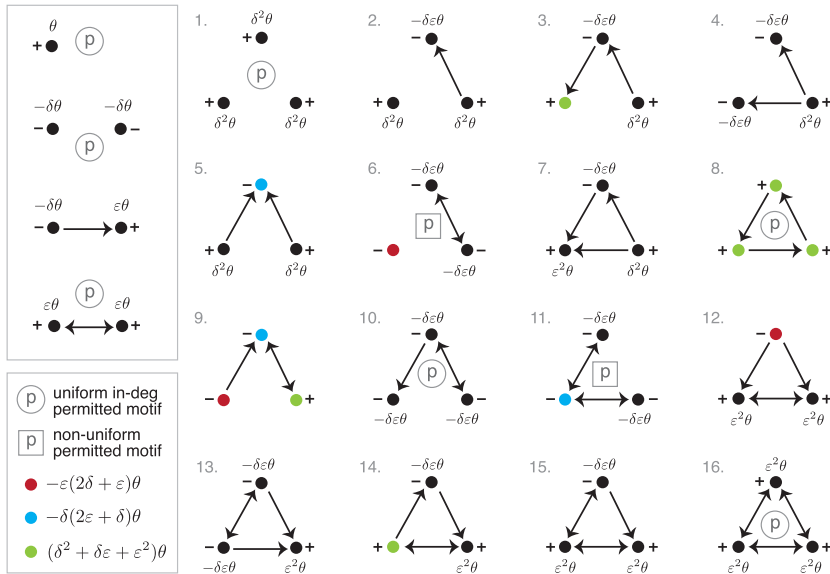


Figure 18: All directed graphs on $n \leq 3$ nodes. (Left) The four non-isomorphic directed graphs with $n \leq 2$. (Right) The 16 non-isomorphic directed graphs with $n = 3$. Each node is labeled with its value of s_i^σ , and its sign, for the full support $\sigma = \{1, \dots, n\}$. Although the s_i^σ values for the colored nodes are non-monomial (bottom left), their signs are also constant throughout the legal parameter range. Thus, using the sign conditions (see Theorem 2), we see that whether a graph is a permitted or forbidden motif is parameter independent for $n \leq 3$. Permitted motifs are labeled with a “p” that lies inside a circle (if the motif is uniform in-degree) or a square (if not). All other graphs are, by definition, forbidden.

CTLN is forbidden, the absence of graphical domination is not sufficient to guarantee that a motif is permitted. Nevertheless, in many cases, graph rules can still be used to determine whether such a motif is permitted or forbidden. In particular, Rule 1 (parity) can be used to determine if a motif is permitted or forbidden provided we know what happens for all proper subsets.

To illustrate this idea, we consider the family of oriented graphs with no sinks on $n \leq 5$ vertices.⁷ Precisely six of these graphs are forbidden motifs despite the absence of graphical domination (see Figure 20). Fortunately, using graph rules, we can identify all permitted motifs with $|\sigma| \leq 4$ that can arise as proper subgraphs in this family, together with their survival rules. Combining this with parity, we will show that the motifs in Figure 20 are all forbidden.

⁷Recall that a directed graph is *oriented* if it has no bidirectional edges.

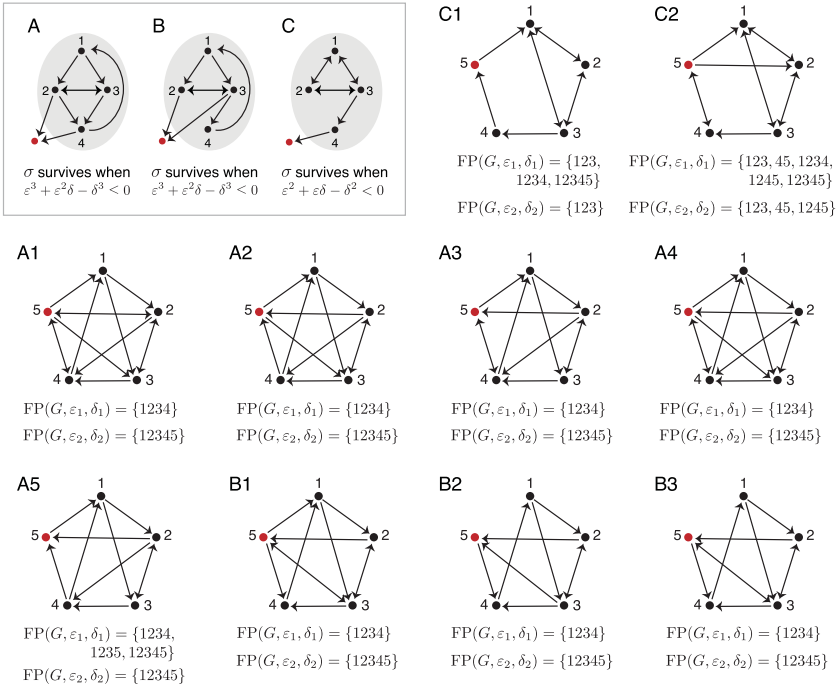


Figure 19: (A–C) The size 4 permitted motifs with parameter dependent survival for the embeddings. (A1–A5) Some example graphs on five nodes with the graph A embedding as a subgraph. As a result $FP(G, \varepsilon, \delta)$ is parameter dependent for these graphs. Below each graph is $FP(G, \varepsilon_1, \delta_1)$ for $(\varepsilon_1, \delta_1) = (0.25, 0.5)$ together with $FP(G, \varepsilon_2, \delta_2)$ for $(\varepsilon_2, \delta_2) = (0.1, 0.12)$. (B1–B3) Some example graphs with the graph B embedding as a subgraph. (C1, C2) Some example graphs with the graph C embedding as a subgraph.




Figure 20: Oriented graphs with no sinks on $n = 5$ that are forbidden motifs despite the absence of inside-in graphical domination.

A.3.1 Survival rules for oriented permitted motifs with $|\sigma| \leq 4$. Figure 21 shows the five permitted motifs with $|\sigma| \leq 4$ that can potentially support fixed points within an oriented graph with no sinks. Note that independent sets of arbitrary size are also permitted motifs, but by Rule 3, an




Figure 21: Permitted motifs of size $|\sigma| \leq 4$ (excluding independent sets) that can arise as subgraphs of an oriented graph with no sinks.

independent set can support a fixed point only if it is a union of sinks, and so no independent set will ever survive in a graph with no sinks.

The first three permitted motifs in Figure 21 are all uniform in-degree 1. Thus, the survival rule for these motifs is given by Rule 2: $\sigma \in \text{FP}(G)$ if and only if no external node $k \notin \sigma$ receives two or more edges from the nodes in σ .

The fourth motif is the disjoint union of a 3-cycle with an isolated node. Recall that a necessary condition for a disjoint union to survive is that each of its components survives (see Proposition 2). Since the isolated node cannot survive in a graph with no sinks, we see that this motif will never survive in our family of $n = 5$ graphs.

Finally, the butterfly graph is permitted by a parity argument. Since the graph has no sinks and no bidirectional edges, there are no fixed point supports of size less than or equal to 2. The only permitted motifs of size 3 are the two 3-cycles—123 and 234—which both survive by Rule 2. Thus, by parity, the full graph must be a fixed point support. The survival rules for the butterfly graph are summarized in the following lemma.

Lemma 21. *Let $\sigma = \{1, 2, 3, 4\}$, and let G be any graph such that $G|_\sigma$ is the butterfly graph in Figure 21. Then $\sigma \in \text{FP}(G)$ if and only if every $k \notin \sigma$ either receives at most one edge from σ , or receives two edges from among nodes 1, 2, and 4.*

Proof. To derive the butterfly graph’s survival rules, let $\sigma = \{1, 2, 3, 4\}$, and consider all possible configurations of outgoing edges from the σ to an added node 5. It is easy to check that if node 5 receives fewer than two edges from σ , there will always be inside-out graphical domination, and so σ will survive. In contrast, if node 5 receives three or more edges from σ , then 5 will outside-in dominate some node in σ , and so σ will not survive.

When node 5 receives exactly two edges from σ , survival is more complicated and depends on which nodes 5 receives from. There are four cases: (i) $1, 3 \rightarrow 5$ (equivalently $3, 4 \rightarrow 5$ by symmetry): 5 outside-in dominates 1, and so σ does not survive; (ii) $2, 3 \rightarrow 5$: 5 outside-in dominates 3, and so σ does not survive; (iii) $1, 4 \rightarrow 5$: 2 inside-out dominates 5, and so σ survives; and (iv) $1, 2 \rightarrow 5$ (equivalently $2, 4 \rightarrow 5$), in which there are no graphical domination relationships, and so we explicitly compute the values

of s_i^σ and check the sign conditions (see Theorem 2). We find that

$$\begin{aligned} s_1^\sigma &= s_4^\sigma = -\delta\theta(\varepsilon^2 + \varepsilon\delta + \delta^2), & s_2^\sigma &= -\delta\theta(2\varepsilon^2 + 2\varepsilon\delta + \delta^2), \\ s_3^\sigma &= -\delta\theta(2\varepsilon^2 + 3\varepsilon\delta + 2\delta^2), \end{aligned}$$

while $s_5^\sigma = \delta\theta(5\varepsilon^3 + 6\varepsilon^2\delta + 5\varepsilon\delta^2 + \delta^3)$. Thus, $\text{sgn } s_5^\sigma = -\text{sgn } s_i^\sigma$ for all $i \in \sigma$, and so σ survives. \square

A.3.2 Proof That the Six Graphs in Figure 20 Are Forbidden. Recall that in an oriented graph with no sinks, no independent set can ever survive as a fixed point support, nor can any subgraph that is a 3-cycle with an isolated node. We can thus restrict attention to the remaining permitted motifs shown in Figure 21.

Figure 20A has 3-cycles 123 and 234, both surviving. The tadpole 1235 does not survive, as node 4 receives two edges from it. There are no restricted subgraphs that are 4-cycles, but we see that 1234 is the butterfly graph, and it survives by Lemma 21. Thus, $\text{FP}(G) \supseteq \{123, 234, 1234\}$, and so the only other possible fixed point support, 12345, is ruled out by parity. Hence graph A is a forbidden motif.

Figure 20B is identical to Figure 20A except for the addition of the $4 \rightarrow 1$ edge. As a result, the 3-cycle 234 no longer survives. In addition, the subgraph on 1234 is no longer the butterfly graph and is not a permitted motif by inside-in domination (node 1 graphically dominates 4). Thus, $\text{FP}(G) \supseteq \{123\}$, and the only other candidate support, 12345, is ruled out by parity. Figure 20B is thus also a forbidden motif.

Similar parity arguments based on subgraph survival can be used for each of the remaining graphs in Figure 20. We leave the remaining graphs as an exercise to readers.

A.4 Associating a Graph to a General Competitive TLN. We have seen that the graph of a CTLN determines many properties of its collection of fixed point supports. Thus, it is natural to consider associating a graph to a general competitive TLN with constant external input $\theta > 0$ and ask if properties of this graph also shape the fixed point supports. Recall that for such networks, Theorems 2 (sign conditions) and 15 (general domination) apply, and either can be used to characterize the fixed point supports of the network.

We associate to each network (W, θ) a directed graph G_W as follows:

$$G_W \text{ has an edge from } j \rightarrow i \text{ (for } i \neq j) \Leftrightarrow W_{ij} > -1.$$

For ease of notation, we will write W_{ij} as $-1 + c_{ij}$ with $c_{ki} < 1$ (to ensure W is competitive). In this case, $j \rightarrow i$ in G_W if and only if $c_{ij} > 0$. Note that G_W is a simple directed graph (no self loops). This generalizes the graph used

in CTLNs. Indeed, if $W = W(G, \varepsilon, \delta)$ is a CTLN network with graph G , then $G_W = G$.

The graph G_W can be used to describe certain aspects of the fixed point supports of the network (see equation 1.1). First, we see that sinks of the graph correspond to WTA (stable) fixed points. Recall that a vertex in a graph is a sink if it has out-degree 0.

Lemma 22. *A TLN (W, θ) has a fixed point supported on a single neuron j if and only if j is a sink of G_W . Moreover, when $\{j\} \in \text{FP}(W)$, the corresponding fixed point is stable.*

Proof. We prove this by computing the values of $s_j^{(j)}$ and $s_k^{(j)}$ for any $k \neq j$, then using Theorem 2 (sign conditions) to determine when $\{j\}$ supports a fixed point. Note that $s_j^{(j)} = \det((I - W_{\{j\}})_j; \theta) = \theta$, and so $\{j\}$ is a permitted motif with positive index. Furthermore, the eigenvalue of $I - W_{\{j\}}$ is θ , which is positive, and so $\{j\}$ supports a stable fixed point in its restricted subnetwork. This fixed point survives as a fixed point of (W, θ) precisely when $\text{sgn} s_k^{(j)} = -\text{sgn} s_j^{(j)} = -1$. Computing $s_k^{(j)}$, we find

$$s_k^{(j)} = \det((I - W_{\{j,k\}})_k; \theta) = \det(I - W_{\{j,k\}}) = \det \begin{bmatrix} 1 & \theta \\ 1 - c_{kj} & \theta \end{bmatrix} = \theta c_{kj},$$

where we have written W_{kj} as $1 - c_{kj}$. Thus, $\text{sgn} s_k^{(j)} = -1$ if and only if $c_{kj} < 0$, which by definition of G_W occurs precisely when $j \not\rightarrow k$ in G_W . By Theorem 2 (sign conditions), $\{j\} \in \text{FP}(W)$ precisely when this condition is satisfied for all $k \neq j$, that is, when j has no outgoing edges, and so j is a sink in G_W . \square

We also find necessary conditions for a fixed point to be supported on exactly two neurons; specifically, we show which subsets of size two are permitted motifs.

Lemma 23. *For a TLN (W, θ) , a pair of neurons $\sigma = \{i, j\}$ is a permitted motif if and only if σ is an independent set or a clique in G_W . Moreover, when $\sigma = \{i, j\} \in \text{FP}(W)$, the corresponding fixed point is stable if and only if σ is a clique.*

Proof. Let $W_{ij} = 1 - c_{ij}$, $W_{ji} = 1 - c_{ji}$, and $\sigma = \{i, j\}$. With this notation, we see $s_i^\sigma = \det((I - W_{\{i,j\}})_i; \theta) = \theta c_{ij}$ and $s_j^\sigma = \theta c_{ji}$. By Theorem 2 (sign conditions), σ is a permitted motif if and only if $\text{sgn} s_i^\sigma = \text{sgn} s_j^\sigma$, that is, $\text{sgn} c_{ij} = \text{sgn} c_{ji}$. These signs agree precisely when (1) $c_{ij}, c_{ji} > 0$ so that $i \leftrightarrow j$ and σ is a clique in G_W or (2) $c_{ij}, c_{ji} < 0$ and σ is an independent set.

Since $|\sigma| = 2$ and $I - W_\sigma$ has positive trace, we see that the eigenvalues of $I - W_\sigma$ are positive (ensuring σ supports a stable fixed point) precisely when $\det(I - W_\sigma) = c_{ij} + c_{ji} + c_{ij}c_{ji} > 0$. Since W is competitive, $c_{ij}, c_{ji} < 1$, and so $\det(I - W_\sigma)$ is positive precisely when $c_{ij}, c_{ji} > 0$, so that σ is a clique. \square




Figure 22: (A) The graph G_W corresponding to the network W in example 8. (B) Plot of the real part of the eigenvalues of $-I + W$ as a function of c , for W in example 9. The matrix (and, hence, the corresponding full support fixed point) is stable throughout the range $0 < c < 0.25$. At $c = 0$, the real and complex eigenvalues have real parts -3 and 0 , respectively.

Lemmas 22 and 23 show that the fixed points of size $|\sigma| \leq 2$ of a general competitive TLN are completely determined by the graph, matching the case of CTLNs. We might then hope that the CTLN survival rules of these fixed points would hold in the general case or that other graph rules would also go through. Example 9 shows that unfortunately, this is not the case.

Example 9 (cautionary example). Consider the competitive TLN (W, θ) with $\theta = 1$ and

$$W = \begin{pmatrix} 0 & -1 + 2c & -1 - 4c \\ -1 + 2c & 0 & -1 - c \\ -1 + 4c & -1 - c & 0 \end{pmatrix}$$

for any c with $0 < c < 0.25$, to ensure W is competitive. This network has corresponding graph $G = G_W$ with $1 \leftrightarrow 2$, $1 \rightarrow 3$, and no other edges (see Figure 22A). In this example, we compute $\text{FP}(W)$ using Lemmas 22 and 23 and Theorem 2 (sign conditions).

By Lemma 22, the only singleton fixed point support is $\{3\}$ since it is the only sink in the graph. By Lemma 23, $\{1, 2\}$ and $\{2, 3\}$ are the only permitted motifs of size 2, but we must check if these survive as fixed points of (W, θ) . For $\sigma = \{1, 2\}$, we have $s_1^\sigma = s_2^\sigma = 2c$ and $s_3^\sigma = 2c^2$; since $\text{sgn } s_3^\sigma = \text{sgn } s_i^\sigma$ for $i \in \sigma$, we see that $\{1, 2\}$ is not a fixed point support, despite being a target-free clique in G_W . In contrast, for $\sigma = \{2, 3\}$, we have $s_2^\sigma = s_3^\sigma = -c$ and $s_1^\sigma = c^2$, satisfying the sign conditions of Theorem 2. Thus, $\{2, 3\}$ is a fixed point support, despite being an independent set that is not a union of sinks of G_W . This shows that both Rules 3 and 4 for CTLNs do not extend to general TLNs.

Finally, consider $\sigma = \{1, 2, 3\}$. We have $s_1^\sigma = c^2$, $s_2^\sigma = 4c^2$, and $s_3^\sigma = 2c^2$. Since the signs all agree, $\sigma \in \text{FP}(W)$. Since σ has uniform in-degree $d = 1$, in any CTLN, σ would also be a fixed point support, but this fixed point

would be unstable by Theorem 5 since $d < \frac{|\sigma|}{2}$. However, for this general TLN (W, θ) , we see that σ supports a stable fixed point for all the allowable values of $0 < c < 0.25$ (see Figure 22B).

Acknowledgments

This work was supported by NIH R01 EB022862 and NSF DMS-1516881. We thank Anda Degeratu for proving an earlier version of Lemma 3, and Chris Langdon for discovering parameter dependence in the graphs of Figure 19.

References

- Curto, C., Degeratu, A., & Itskov, V. (2012). Flexible memory networks. *Bull. Math. Biol.*, *74*(3), 590–614.
- Curto, C., Degeratu, A., & Itskov, V. (2013). Encoding binary neural codes in networks of threshold-linear neurons. *Neural Comput.*, *25*, 2858–2903.
- Curto, C., & Morrison, K. (2016). Pattern completion in symmetric threshold-linear networks. *Neural Computation*, *28*, 2825–2852.
- Hahnloser, R. H., Sarpeshkar, R., Mahowald, M., Douglas, R., & Seung, H. (2000). Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, *405*, 947–951.
- Hahnloser, R. H., Seung, H., & Slotine, J. (2003). Permitted and forbidden sets in symmetric threshold-linear networks. *Neural Comput.*, *15*(3), 621–638.
- Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.*, *79*(8), 2554–2558.
- Morrison, K., & Curto, C. (2018). Predicting neural network dynamics via graphical analysis. In R. Robeva & M. Macaulay (Eds.), *Algebraic and combinatorial computational biology*. Amsterdam: Elsevier.
- Morrison, K., Degeratu, A., Itskov, V., & Curto, C. (2016). *Diversity of emergent dynamics in competitive threshold-linear networks: A preliminary report*. <https://arxiv.org/abs/1605.04463>
- Morrison, K., Langdon, C., Parmelee, C., Geneson, J., Degeratu, A., Itskov, V., & Curto, C. (2019). *Dynamically relevant motifs in inhibition-dominated threshold-linear networks*. Manuscript in preparation.
- Xie, X., Hahnloser, R. H., & Seung, H. (2002). Selectively grouping neurons in recurrent networks of lateral inhibition. *Neural Comput.*, *14*, 2627–2646.

Received April 3, 2018; accepted September 11, 2018.