Symmetries in Network Games

Sarah Taylor*

November 6, 2025

Abstract

Many economic settings are network games, where who agents interact with shapes their behaviour. Characterising equilibrium behaviour in these games is challenging because incentives vary with network position and the space of possible networks is large. We study when identical network incentives lead agents to take the same equilibrium actions or receive the same treatment from a social planner. We introduce algebraic graph theory tools that leverage network symmetries, which fold up the network so that nodes on either side of the fold occupy identical network positions. In unique or extremal equilibria, agents with the same network position must take the same action. These equilibria can admit tractible comparative statics. For interventions, symmetries structure how changes in incentives flow through the network, revealing a contrast in large budget settings. With complements, positionally identical agents are targeted the same. With substitutes, targeting can differentiate among equivalent agents by redistributing their original total equilibrium action amongst them.

*Jesus College, University of Cambridge (email: srt44@cam.ac.uk). I am grateful to Matt Elliott for his insightful comments and supervision. I also thank Sanjeev Goyal, Ben Golub, Christian Ghiglino, and Hamid Sabourian for their helpful comments on this work. Any remaining errors are the sole responsibility of the author. This work was supported by the Economic and Social Research Council [award reference ES/J500033/1].

1 Introduction

A central question in economics is why rational agents achieve different outcomes. The standard answer is that these agents have heterogeneous incentives. In many economic settings, this heterogeneity appears as differences in who we interact with. It shapes people's voting decisions, firms' supplier choices, and refugees' search for work upon arriving in a new country.[1] These are all examples of *network games*: settings where agents care not only about their own actions but also about those of their neighbours.

Systematically modelling incentive heterogeneity in non-networked games is usually not difficult. Agents either have explicitly different parameters in their utility functions or differing functional forms. However, controlling heterogeneity is more challenging in network games. Agents' incentives vary with the network's structure, but the space of distinct network configurations is typically large. This paper asks when will identical network incentives imply identical equilibrium behaviour, and when interventions should be equal or differentiated?

We begin with a simple organising idea: in network games, agents in the same network position face the same network incentives.[2] To think systematically about when two agents occupy the same network position we utilise a mathematical tool called a *network symmetry*.[3] In a star network, for example, all peripheral nodes are equivalent in terms of their network position: each links only to the central hub, which has many neighbours (see Figure 1).

Network symmetries create a framework which organises incentive heterogeneity in network games. Using this framework, we determine sufficient conditions for when players in identical network positions exhibit identical equilibrium behaviour and when these equilibria admit tractable comparative statics. Our approach also reveals that identical incentives do not always imply identical treatment in targeting interventions. Strategic complements require equal treatment, whilst substitutes can necessitate differentiated treatment.

The existing networks literature has largely grappled with incentive heterogeneity in two ways: by assuming the full network structure is not strategically relevant or, when possible, solving explicitly for equilibria. The former is usually done by assuming agents can only observe some limited local information such as their degree (Galeotti et al. [2010]) or neighbourhood of nodes a fixed distance from them (Chaudhuri et al. [2024]), or by replacing the realised network with its limiting distribution (Jackson and Yariv [2007], Golub and Jackson [2010], Acemoglu et al. [2015]). These approaches yield analytical traction, but run into problems when networks with the same local features can produce different actions

[1]Bond et al. [2012] shows in an enormous Facebook RCT that an "I Voted" message from close friends increased political self-expression, information seeking, and real-world turnout. Chaney [2014] uses data on French firms and finds they expand through contact networks. Beaman [2012] finds that larger co-ethnic networks improve political refugees' employment outcomes via information/referrals.

[2]In network games agents have some individual incentive to take an action (you buy a t-shirt in favourite colour) as well as a network incentive to take an action (you buy a t-shirt because your friends own clothes from the same brand). To ensure that the only source of variation in network incentives is the network structure, network games usually assume that network incentives feature in agents' utility functions using the same parameters.

[3]These are called graph automorphisms in mathematics texts.

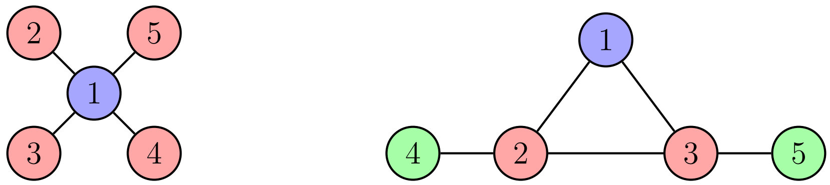

Figure 1: In the star network (left) nodes $\{2,3,4,5\}$ occupy the same network position. In the core-periphery network (right) $\{2,3\}$ share a position, as do $\{4,5\}$.

owing to different non-local features such as differing maximum path lengths or the presence of cycles.

The latter approach usually ranks agents by various centrality measures (which utilise the full network) to solve for equilibria and prescribe welfare-improving interventions (Ballester et al. [2006];Galeotti et al. [2020],Akbarpour et al. [2025]). These measures provide precise characterisations of equilibrium actions, but are somewhat opaque in that they hard to interpret analytically in terms of network architecture, and do not usually apply outside of settings with strategic complements or linear best responses. Existence and uniqueness can still be obtained under certain conditions in more general settings, but often cannot offer a description of the equilibrium structure (Bramoullé et al. [2014]; Allouch [2015],Parise and Ozdaglar [2023],Zenou and Zhou [2024]).

Our symmetries based framework complements both of these approaches. It gives us a mathematical language we can use to flex the network's heterogeneity without discarding its full structure. Moreover, by formalising the notion of positional equivalence, we can also use some powerful but previously un-utilised results from algebraic graph theory to analyse the structure of the centrality measures typically used in network games. This enhances the interpretability of these centrality-based results. We demonstrate the strength of our approach in two applications.

Our first application builds on the idea long-standing idea that symmetry simplifies economic analysis. Classic studies focus on symmetric equilibria, where relabelling players does not affect payoffs (Nash [1951], Dasgupta and Maskin [1986], Harsanyi and Selten [1988], Milgrom and Roberts [1994]; see Plan [2023] for a recent extension). In conventional symmetric games, agents have the same feasible actions and receive the same payoffs when facing the same action profiles. Network games, by design, do not meet these conditions because agents' positions differ, causing differing incentives. We therefore use a relaxed version of the symmetric equilibrium, where agents in the same network position choose the same action. This is only generally possible in games where agents in identical positions have the same non-network incentives. Beyond this, we remain general and do not restrict agents' utility functions.

We leverage symmetries to classify equilibrium structures and derive comparative statics in settings where the existing literature struggles to find traction. Firstly, we determine when identical network incentives imply identical equilibrium outcomes. We find that this occurs when the equilibrium is unique, or is the extremal equilibrium of a game of strategic complements. This allows us to derive equilibrium structure in settings where information is scarce because all we have is some sort of existence result. Unlike arbitrary equilibria in network games, these position-invariant equilibria admit symmetry-based comparative statics. We identify sufficient limitations on link structure that permit comparative statics in settings such as non-linear best responses with strategic substitutes. Our network symmetries framework also suggests a new type of comparative static that simplifies analysis through explicitly controlling heterogeneity: compare networks with the same number of distinct network positions, rather than the same number of nodes.

Our second application uses network symmetries to classify optimal targeting in networks. We study an incentive design problem, where a planner adjusts agents' equilibrium actions by altering their non-network returns, subject to a budget. Unlike the previous application, we allow agents in the same network position to have different non-network incentives. We focus on the widely used quadratic network game, which has been used in settings such as price competition, peer effects in education, and technology adoption.[4]

The solution to this incentive design problem at first appears complicated, as it involves solving a high-dimensional system of interdependent first-order conditions. However, Galeotti et al. [2020] shows that it can be cleanly formulated in terms of the eigenvalues and eigenvectors of the matrix representing the network's links (the *adjacency matrix*). They find that in high-budget settings, the optimal intervention converges to a single eigenspace. When the game exhibits complements, it converges to the eigenspace associated with the largest eigenvalue, whilst substitutes converge to the eigenspace associated with the smallest eigenvalue.

Our symmetries framework builds on Galeotti et al. [2020]'s results. We demonstrate that the planner may not always want to treat agents in the same network position the same way. This is despite these agents having identical network incentives, and thus an identical impact in terms of changing the actions of others in the network. We find that when there are complements, the planner will always want to treat agents with same network position in the same way, but this is not always the case for substitutes. Here, the planner chooses an intervention that reallocates equilibrium action between agents in the network position, whilst keeping their total equilibrium action fixed.

We also expose a paradox for settings even when the planner's budget is not large: more symmetry implies that identical-looking agents are treated less equally. Specifically, networks with larger numbers of agents that occupy the same network position (and thus less heterogeneity in terms of their network incentives) have optimal interventions wherein a larger portion of the budget is spent on redistributing action between agents, relative to those with smaller groups of position -equivalent agents.

These conclusions are possible because we can use a deep mathematical relationship that exists between the eigenvectors and eigenvalues of a network and its symmetries.[5] The symmetries partition the eigenvectors associated with the adjacency matrix based on how they act on agents in the same network position. One set is constant, and the other sums to

[4]Price competition: Candogan et al. [2012], Fainmesser and Galeotti [2016]. Peer effects: Calvó-Armengol et al. [2009]. Technology adoption: Banerjee et al. [2013]

[5]Specifically, an implication of Maschke's theorem.

zero. This is useful because the total intervention can be expressed as a weighted sum of the adjacency matrix's eigenvectors (as shown in Galeotti et al. [2020]). This partition yields a clean economic interpretation. The eigenvectors that are constant on each set of position-equivalent agents contribute to the total intervention by shifting these agents' actions up or down by the same linear amount. Thus, these eigenvectors roughly capture pure spending interventions. Conversely, the eigenvectors that sum to zero contribute to the intervention by simply shifting the total equilibrium action around between agents in the same position. These capture balanced-budget reallocations that are implementable through subsidies and taxes.[6]

Our symmetries-based characterisation of the entire eigenspace provides a set of methods that apply to the many works that use eigenspaces in a network setting, and thus extends beyond Galeotti et al. [2020]. For example, Liu and Tsyvinski [2024] examines shock propagation in US production networks, and Banerjee et al. [2013] ranks nodes' information-spreading capacity using a measure akin to eigenvector centrality. To the best of our knowledge, we are the first to use this result in the economics network literature.

Our work formalises the role of network symmetries and links them to mathematical group theory. Whilst formal treatments are network symmetries are limited, two notable exceptions are recent works that apply them in games with linear-quadratic utilities with strategic complements. Allouch and Bhattacharya [2025] introduces a "key class" partition, related to the partition of agents by network position, to identify a group of nodes that a planner should remove to reduce aggregate activity. Chaudhuri et al. [2024] analyse games with incomplete network information, showing that players with isomorphic local neighbourhoods and beliefs choose identical actions. Our targeting application generalises these by characterising the entire eigenspace. We also contribute to work on how equilibrium actions align with eigenvectors, notably Bonacich and eigenvector centralities, which emerge from linear manipulations of the adjacency matrix (Ballester et al. [2006], Bramoullé et al. [2014]).

The remainder of the paper proceeds as follows. Section 2 introduces the mathematical formalities of network symmetries and orbits. Section 3 presents our first application of characterising equilibria of symmetric network games and introducing comparative statics. Section 4 applies symmetries to targeting problems and shows how they structure optimal interventions. Section 5 concludes.

2 Symmetry Groups of Networks

Algebraic graph theory studies the relationship between various algebraic objects and network properties. It has three main branches: linear algebra, group theory, and graph invariants.[7] The first is widely used in network economics. We focus on the second: a rich mathematical literature that leverages groups to learn more about network structure. Groups are one of

[6]We note that Galeotti et al. [2024] focuses on using spectral methods to speak to a balanced budget intervention. Their approach involves a per-unit tax balanced at the equilibrium level of activities, whereas our observation suggests balancing via a costly action that alters marginal returns.

[7]These invariants typically pertain to polynomials e.g. the chromatic polynomial, which counts colourings.

the standard mathematical tools for dealing with symmetries of various kinds.[8] We introduce some key mathematical definitions below. The precise definition of a group is not important for this work, however, we include it in the appendix.

Throughout this text, we consider unweighted and undirected networks. A network Γ consists of a pair $(V(\Gamma), E(\Gamma))$. $V(\Gamma)$ is the set of vertices (or nodes) which we will denote $N = 1, \dots, n$, and $E(\Gamma)$ is the set of edges (links) which are unordered pairs $\{i, j\}$ with in $i, j \in N$. We may represent Γ by a symmetric binary matrix, with a one in entry ij if there is an edge between ij and zero otherwise. Because we often think about the network and its adjacency matrix interchangeable, we abuse notation and also call the adjacency matrix Γ.

A *permutation*, Π, on a finite set S is a bijection from S onto itself. Four our purposes S be the set of nodes of Γ, which we denote as $V(\Gamma)$ when it is not clear from context.

Definition 1 (Network Symmetry). *A* symmetry *of an undirected network Γ is a permutation $\Pi : V(\Gamma) \to V(\Gamma)$ such that $\{i, j\}$ is an edge in Γ if and only if $\{\Pi(i), \Pi(j)\}$ is an edge in Γ. The set of all symmetries is called the* symmetry group of Γ, *which we shall denote by G.*

Figure 2 demonstrates an example of a network symmetry.

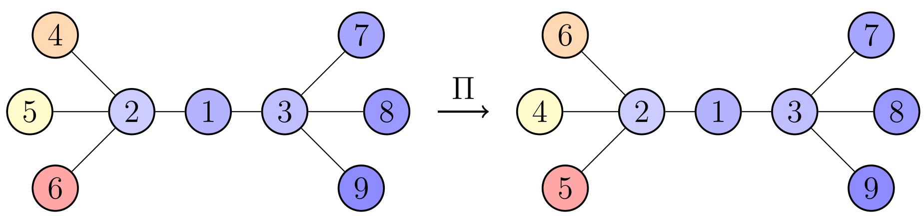

Figure 2: A network symmetry where the permutation Π maps $4 \mapsto 5$, $5 \mapsto 6$, $6 \mapsto 4$, and leaves all other nodes fixed.

It turns out that we can use the symmetry group of Γ to partition $V(\Gamma)$ into disjoint sets of symmetrically equivalent nodes, as described in the remark below.

Definition 2 (Symmetric Equivalence). *We say $i \sim j$ if there exists a symmetry Π such that $\Pi(i) = j$.*

This defines an *equivalence relation* on the nodes of Γ. To see this, notice that for reflexivity, the identity operation is a network symmetry, and so every node is similar to itself. For symmetry, notice that since permutations are bijections, the inverse mapping Π^{-1} exists and also preserves adjacency; that is, if an edge exists between $\Pi(h)$ and $\Pi(k)$, then applying Π^{-1} shows that the corresponding edge exists between h and k.[9]. Transitivity follows from the fact that the composition of two symmetries is again a symmetry.[10] Every equivalence relation partitions the set on which it acts into disjoint sets called equivalence classes. These will be important for our analysis going forward:

[8]Our section on targeting in fact relies on a surprising connection between linear algebra and group theory.

[9]Explicitly: $\{h, k\} \in E(\Gamma) \iff \{\Pi(h), \Pi(k)\} = \{h', k'\} \in E(\Gamma)$ then the condition $\{h', k'\} \in E(\Gamma) \iff \{\Pi^{-}(h'), \Pi^{-}(k')\} = \{h, k\}$ is automatically satisfied for all edges

[10]These properties all actually hinge on the fact that G is a group.

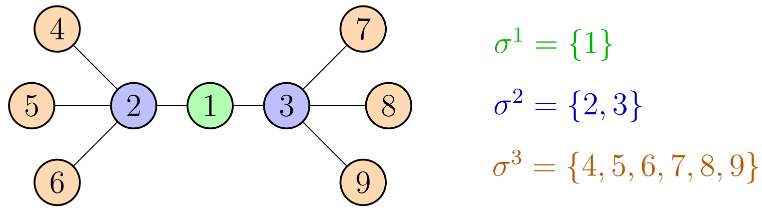

Figure 3: The orbits of the symmetry group partition the nodes into geometrically equivalent sets.

Definition 3 (Orbits)**.** *Let* $\{\sigma^1, \ldots, \sigma^r\}$ *be the equivalence classes induced on* $V(\Gamma)$ *by* $\sim$*. We call* $\{\sigma^1, \ldots, \sigma^r\}$ *the* orbits *of* $V(\Gamma)$ *under* G*.*

Figure 3 shows our running example partitioned into orbits.

Informally, if two nodes are in the same orbit then Γ 'looks the same' to all these nodes, in the sense that at every distance from them their neighbourhoods are identical: there is a direct mapping between their neighbours and their neighbours' neighbours etc.. Orbits are useful because they collect together geometrically equivalent nodes, and many network games are solved using geometric properties of the nodes (for example, path centrality). We may, in a sense, use any agent in an orbit as a representative of all of its peers.

Now, if i and j are in the same orbit, and we define i's neighbours in a particular orbit to be $N_i^k = N_i \cap \sigma^k$, then $|N_i^k| = |N_j^k|$. This is because if $h \in \sigma^k \cap N_i$ then $\Pi(h) \in N(\Pi(i)) = N_j$, and clearly $\Pi(h) \in \sigma^k$ if $h \in \sigma^k$. Thus we can use these orbits to reduce Γ to a completely asymmetric version of itself, called the *quotient graph* and denoted Γ/G, with adjacency matrix $(\Gamma/G)_{mk} = N_{mk}$. The quotient graph is weighted (by the positive integers) and directed. Figure 4 demonstrates an example of it below.

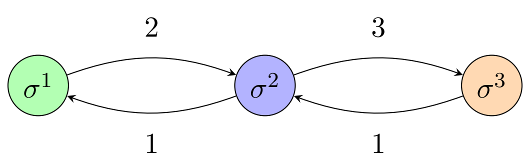

Figure 4: The quotient graph of our running example, where each node now represents an orbit.

The quotient graph is useful for several reasons, including that its spectrum gives us useful information about the spectrum of Γ, which we shall explore further in Section 4.

3 Equilibrium Structure and Comparative Statics

We study a widely-used class of network games where players' payoffs depend entirely on their own actions and those of their neighbours. In these games, players in the same network position (i.e. orbit) have identical strategic incentives. These are sometimes called *symmetric network games*. The section below analyses these games. We examine when nodes in the same position take the same action in equilibrium. We call these *orbit-invariant* equilibria and identify sufficient conditions for their existence. Orbit-invariant equilibria enable tractable comparative statics, and we demonstrate their applications.

3.1 Model: Symmetric Network Games

We consider a set of agents $N = \{1, \ldots, n\}$, and let $\Gamma(n)$ denote the set of unweighted and undirected networks on n nodes. Let X be a common strategy set available to all agents, so that the full strategy space is X^n. Each agent $i \in N$ selects an action $x_i \in X$, and together these form a strategy profile $\mathbf{x} = (x_1, \ldots, x_n) \in X^n$. Agent i's payoff depends on the strategy profile and the underlying network structure, and is given by the utility function $u_i : X^n \times \Gamma(n) \to \mathbb{R}$. For convenience, we shall often pick out i's strategy and denote strategy profiles by the tuple (x_i, x_{-i}), where $x_{-i} = (x_j)_{j \neq i}$.

Let $N_i(\Gamma)$ denote the set of neighbours of agent i in the network $\Gamma \in \Gamma(n)$. We impose two natural restrictions on the utility functions to ensure that only agents' network position determines their incentives and payoffs:

Assumption 1 (Unordered dependence on neighbours)**.** *An agent's payoff depends only on their own action and the multiset of the actions taken by their immediate neighbours. Formally, for any strategy profiles $\mathbf{x}, \mathbf{x}' \in X^n$, and a fixed network Γ, we have*

$$u_i(x_i, x_{-i}; \Gamma) = u_i(x'_i, x'_{-i}; \Gamma)$$

whenever $x_i = x'_i$ and $\{x_j\}_{j \in N_i(\Gamma)} = \{x'_j\}_{j \in N_i}$.[11]

Assumption 1 implies that agents are agnostic as to which neighbour takes which action. Instead, only the collection of actions matters. This assumption is satisfied in many network games. It is for example satisfied in the many models where payoffs depend on the sum of neighbours' actions.

Assumption 2 (Agent Anonymity)**.** *If two agents i, j have the same number of neighbours, take an identical action, and their neighbours take the same multiset of actions, then i and j receive the same payoff. That is,*

$$u_i(x_i, x_{-i}; \Gamma) = u_j(x'_j, x'_{-j}; \Gamma)$$

whenever $x_i = x'_j$, $|N_i(\Gamma)| = |N_j(\Gamma)|$ and i and j's neighbours actions take the same set of actions: $\{x_k\}_{k \in N_i} = \{x'_k\}_{k \in N_j}$.

[11] Here $\{x_j\}_{j \in N_i(\Gamma)}$ denotes the *multiset* of neighbours actions, so that repeated elements are allowed.

Assumption 2 states that only an agent's network position determines their payoffs, and thus strategic incentives are fully dictated by network position. It means, for instance, that two agents in the same orbit must have the same parameters in their utility functions.[12] Assumption 2 is widely used in the networks literature when authors want the network to be the only source of heterogeneity for agents, although it is often relaxed for applied works. We drop Assumption 2 in our model in Section 4.

Strategic symmetries. A network symmetry is essentially a relabelling of the network which leaves neighbours unchanged. However, economists have long had notion of a strategic symmetries in general (non-network) games: a relabelling of the agents in the game which leaves payoffs unchanged. We call these symmetries *strategic symmetries*, to differentiate them from network symmetries. To formalise the notion of a strategic symmetry, we begin with a classical definition due to von Neumann and Morgenstern.

Definition 4 (von Neumann and Morgenstern (1953))**.** *A game* $(N, X^n, \mathbf{u})$ *is said to be* symmetric *with respect to a permutation* $\Pi : N \to N$ *if* $X_i = X_{\Pi(i)}$ *for all* $i \in N$*, and the permutation preserves the payoff structure:*

$$u_{\Pi(i)}(x_1, \dots, x_n) = u_i(x_{\Pi(1)}, \dots, x_{\Pi(n)}).$$

It is important to distinguish this from the stronger notion of a *totally symmetric game.*[13] In a totally symmetric game, *every* permutation of players is a symmetry of the game. In other words, total symmetry requires that any set of players can be swapped and the game remains unchanged. By contrast, a game with symmetry (but not total symmetry) may admit only a subgroup of permutations as symmetries.

Many games familiar to economists exhibit total symmetry. For example, in the standard Cournot and Bertrand oligopoly models, permuting any set firms yields leaves the game unchanged. However, network games are generally not totally symmetric by design: the fact that agents are heterogenous in who they interact with is what makes the role of the network salient.

3.2 Orbit-Invariant Equilibria: Existence and Comparative Statics

We now define a natural network game analogue to the totally-symmetric equilibrium, wherein *all* agents take the same action. We now focus on equilibria where all agents in the same network position take take the same action. We call these equilibria *orbit-invariant* equilibria:

[12]For example, in the standard quadratic linear utility function $u_i(x_i, x_{-i}) = x_i(b_i + \beta \sum_{j \in N_i} x_j) - c_i x_i^2$, agents in the same orbit must have the same value for $b_i \in \mathbb{R}$ and $c_i \in \mathbb{R}^+$. β typically lies in $[-1, 1]$, and is usually assumed to be the same for all agents in network games.

[13]Plan [2023] formalises this distinction and provides a thorough discussion of its implications.

Definition 5 (Orbit-Invariance)**.** *Let* Γ *be a graph on* n *nodes with orbits* $\sigma^1, \ldots, \sigma^r$ *under its symmetry group* G*. We say that a tuple* $(x_1, \ldots, x_n)$ *is orbit-invariant if* $i, j \in \sigma^k$ *implies that* $x_i = x_j$*.*

Thus, for an action profile to be orbit-invariant all nodes in the same orbit must take the same action. In some ways, these profiles reflect a natural form of behavioural consistency: players in the same network position behave identically. In the section below, we study when such orbit-invariant equilibria arise. We then study they whether inherit some of the tractable comparative statics of their totally-symmetric counterparts, and also suggest a new type orbit-based comparative statics.

3.2.1 Games with Orbit-Invariant Equilibria

Here, we explore the relationship between the symmetry structure of the network and the strategic symmetries of its associated games. We know from the definitions section above that if two agents occupy nodes which are in the same orbit under the symmetry group of the network, then they occupy the same structural positions in the network. Intuitively, if two agents are structurally indistinguishable and their utility functions depend only on their own and their neighbours' actions, then they must be strategically interchangeable. This leads to the following observation:

Lemma 1. *Every symmetry of the network* Γ *is a symmetry of the game* $(N, X^n, \mathbf{u}(\mathbf{x}, \Gamma))$*.*[14]

The proof of the above is fairly intuitive. Suppose that Π is a network symmetry and $\mathbf{x}$ is a particular strategy profile. If we apply Π to $\mathbf{x}$ and afterwards node j plays node i's strategy, then all of j's neighbours are now playing the strategy originally played by i's neighbours, and so by Assumption 2, j now gets i's payoff. A formal proof of this intuition can be found in the appendix.

This result is an important starting point for our analysis because it links the combinatorial structure of the network and the equilibrium behaviour it supports. Symmetries of games are important because they can often tell us about the structure of a wide class of equilibria for which explicit solutions are difficult to find.

Proposition 1 (Orbit-invariant equilibria in pure strategies)**.** *Suppose that* i *and* j *are in the same orbit of* Γ *under* G*. Then:*

- *If the game* $(N, X^n, \mathbf{u}(\mathbf{x}, \Gamma))$ *has a* unique *equilibrium* $\mathbf{x}^*$*, then* $x_i^* = x_j^*$*; and*
- *If the game* $(N, X^n, \mathbf{u}(\mathbf{x}, \Gamma))$ *is* strictly supermodular *with maximal equilibrium* $\overline{\mathbf{x}}^*$ *and minimal equilibrium* $\underline{\mathbf{x}}^*$*, then* $\overline{x}_i^* = \overline{x}_j^*$*, and* $\underline{x}_i^* = \underline{x}_j^*$*.*

The following also immediately follows from Lemma 1 once we apply the following classic result due to Nash [1951], which states that any finite game has a symmetry-invariant equilibrium in mixed strategies.[15]

[14] For those familiar with group theory: G is a subgroup of $\mathcal{G}$

[15] A game is said to be finite if it has a fixed number of players and a finite action space.

Corollary 1. *Suppose X is finite. Then the game $(N, X^n, \mathbf{u}(x, \Gamma))$ always has an equilibrium in mixed strategies wherein all nodes in the same orbit of Γ play the same strategy.*

These results are helpful in the many settings where we cannot solve for equilibria explicitly, but know they satisfy certain properties. They are especially relevant for existing equilibrium uniqueness results for general network games.[16] However, relating these unique equilibria to the network's structural properties is usually difficult in these very general settings. Similarly, many network games of interest, such as diffusion models, are strictly supermodular but lack explicit equilibrium solutions. Proposition 1 complements these results by adding a structural descriptor (orbit-invariance).

Our results are also useful because symmetries can simplify the computation of equilibria. This is important as network games often involve high-dimensional solutions. Restricting the equilibrium search to orbit-invariant strategy profiles can reduce computational complexity.

The role of symmetries in describing equilibria is robust to changes in agents' payoff functions and strategy spaces. This contrasts with network centrality measures, which provide fine-grained predictions about equilibrium behaviour but each new economic setting requires a new centrality measure. For example, Bonacich centrality determines which agents take a higher equilibrium action in the quadratic linear utility model (Ballester et al. [2006]). But identifying which agents are more influential in diffusing ideas and innovations relies on eigenvector centrality (Golub and Jackson [2010], Banerjee et al. [2013]).[17] Network symmetries provide a different descriptive lens. They apply widely and identify agents who take the same action in equilibrium. However, they cannot describe how agents' actions differ across orbits.

Network centrality measures are also usually defined recursively in terms of the network's adjacency matrix. These rarely make the structure of ties and rankings transparent. Thus, *a priori* comparisons across different structural positions typically require computing the index. Symmetries, on the other hand, are easily identifiable by inspection and immediately imply orbit-constancy.

Our focus on orbit-invariant equilibria can also be considered compatible with centrality-based analysis. A notable contribution that generalises centrality measures is Sadler [2022]. This work introduces *ordinal centrality*, a general measure that ranks agents based on an iterative comparison of their neighbourhood sizes. Ordinal centrality ranks equilibrium actions in general network games of strategic complements. A key assumption of ordinal centrality is that it is orbit-invariant. Sadler [2022] shows that several important centrality measures are in fact ordinal centralities - and thus orbit-invariant.

Our ability to map symmetries into matrices also gives us an easy, algebraic way to prove that the symmetries of Γ are also the symmetries of Bonacich centrality, when all nodes have the same utility function:

---

[16] See, for example: Zenou and Zhou [2024].

[17] Bonacich centrality ranks agents by counting the number of walks emanating from them. Eigenvector centrality instead ranks nodes by their entries in the eigenvector associated with the largest eigenvalue of the network's adjacency matrix.

Remark 1. *Let* Bonacich centrality *of an undirected network* Γ *be defined by* $\mathbf{c}^B = b(\mathbb{I} - \beta\boldsymbol{\Gamma})^{-1}\boldsymbol{\Gamma}\mathbf{1}$*, where* β *and* $b \in \mathbb{R}$*, and* $\mathbf{1}$ *is the length* n *vector of ones. Then, if* i*,* j *are in the same orbit of the symmetry group of* Γ*,* $c_i^B = c_j^B$*.*

The proof may be found in the appendix, but relies on a standard fact that $\Pi\Gamma = \Gamma\Pi$.

We have so far shown that every symmetry of the network must be a strategic symmetry of the game. We now derive a sufficient condition for when every strategic symmetry of the game must be a strategic symmetry of the network.

Proposition 2. *Suppose that* u *is additively separable in the following sense:* $\exists\ \nu : X \to \mathbb{R}$*,* $\mu : X^2 \to \mathbb{R}$ *and injection* $f : \mathbb{R} \to \mathbb{R}$ *such that*

$$u(x_i, x_{-i}, \Gamma) = f\left(\nu(x_i) + \sum \Gamma_{ij}\mu(x_i, x_j)\right).$$

Moreover, assume that $\exists(\tilde{x}, x^0) \in X^2$ *such that* $\mu(\tilde{x}, x^0) = 0$ *and* $\exists x^1 \in X$ *such that* $\mu(\tilde{x}, x^1) \neq 0$*.*

Then Π *is a symmetry of the game* $\iff$ *it is a symmetry of* Γ*.*

The requirements on the existence of $(\tilde{x}, x^0)$ and $(\tilde{x}, x^1)$ may at first appear quite restrictive. However, because this is a sufficient condition, it may be that the actual requirements are somewhat weaker. The reason for these precise constraints is that we need to construct an action profile where only a given node's neighbours take an action which give a non-zero payoff. Moreover, many network games do satisfy this additively separable property: for example threshold games and congestion games. This converse result—that every symmetry of the game is also a symmetry of the network— implies that all strategic equivalences among players arise solely from the network topology, rather than from exogenous invariances in the payoff structure.

3.2.2 Comparative Statics for Orbit-Invariant Equilibria

The section above shows that orbit-invariant equilibria occur in variety of settings. We now analyse their comparative statics. Motivated by studying settings with strategic substitutes or non-linear best responses, we examine when orbit-invariant equilibria inherit the tractable comparative statics of totally symmetric games. We then propose a new kind of orbit-based comparative static.

Robust comparative statics on parameters on one–factor networks. Milgrom and Roberts [1994] derive conditions under which totally-symmetric games admit symmetric equilibria that inherit the same monotone comparative statics on parameters as the game's best response function.[18] This result is useful as it studies a very general setting: it requires only mild continuity, and makes no assumptions on the sign or curvature of the best response function's derivative. We wish to extend this useful result to orbit-invariant equilibria.

[18]Theorem 1 of Milgrom and Roberts [1994].

As in Milgrom and Roberts [1994], we consider the parameter-space of the game to be a partially ordered set, Θ. We also make their key assumption which forces agents' best responses to cross the axis at least once.

Assumption 3 (MR94)**.** *Let $f(x,\theta) : [0,1] \times \Theta \to [0,1]$, where Θ is a partially ordered set. We say f satisfies MR94 if:*

- *$f(0,\theta) \geq 0$ and $f(1,\theta) \leq 0$, and for each θ, and*
- *$f(\cdot,\theta)$ is continuous but for upward jumps: if for all y, $\limsup_{x\to y^-} f(x,\theta) \leq f(y,\theta) \leq \liminf_{x\to y^+} f(x,\theta)$.*

The crux of Milgrom and Roberts [1994]'s proof for totally symmetric games is to write equilibrium conditions as a fixed-point equation in a single scalar. In other words, if $\phi(x,\theta)$ is the best response function, set $f(x,\theta) = \phi(x,\theta) - x$, and find the roots of $f(x,\theta)$. In one dimension a parameter shift moves the graph of the scalar map in a single direction, yielding monotone changes in its fixed points. However, the arbitrary heterogeneity of network games typically prevents us from doing this - instead we have to solve for a system of fixed points. In many dimensions the parameter does not induce a total order in $\mathbb{R}^n$, so fixed points can change non-monotonically. To collapse the role of the network in the best response function to a single scalar, we restrict our attention to networks whose orbits display a particular kind of structure. We call such networks *one-factor networks*:

Definition 6 (one-factor networks)**.** *We say that a network Γ is a one-factor network if the adjacency matrix of its quotient network, Γ/G, has rank one.*

One-factor networks are a family of networks that include stochastic block models whose matrix of edge probabilities is rank one, hierarchical networks consisting wherein each layer is a regular network, and vertex-transitive networks. In a one-factor network, each orbit connects to every other orbit in the same ratio. For example, if half of one orbit's are with orbit r, then half of any other orbit's must also be with orbit r. This is particularly helpful if the network game is *additive* in the following sense:

Definition 7. *Suppose $X \subseteq \mathbb{R}$. We say a network game is* additive *if each player i's best response function can be written as $\phi_i(\sum \Gamma_{ij} x_j; \theta)$, where $x_i = \phi_i : \mathbb{R} \times \Theta \to \mathbb{R}$.*

Our extension to orbit-invariant equilibria relies on the idea that we can solve for them on the quotient graph. We solve one best response function per orbit. Instead of taking $\sum \Gamma_{ij} x_j$ as an input, it now takes $\sum (\Gamma/G)_{rs} \tilde{x}_s$, where $\tilde{x} \in [0,1]^R$ is the vector representing each orbit's action. Recall that an $m \times n$ matrix M has rank one if there exist $b, w \in \mathbb{R}^m$ such that $M = b\, w^\top$. Thus, for nodes in orbit σ^r we have that

$$\sum_s (\Gamma/G)_{rs}\, \tilde{x}_s = \sum_s \left(b\, w^\top\right)_{rs} \tilde{x}_s = b_r \sum_s w_s \tilde{x}_s,$$

and so the role of the network in best responses collapses to a single sufficient statistic. We may then apply Milgrom and Roberts [1994] to obtain the following result:

Proposition 3 (Comparative statics for one-factor networks)**.** *Let Γ be a one-factor network. Consider an additive network game with strategy space $X = [0, 1]$. If each $\phi_i(\cdot; \theta)$ satisfies Assumption 3, then there exist a lowest and a highest orbit-invariant equilibrium, $\underline{x}_i(\theta)$ and $\overline{x}_i(\theta)$. Moreover, if each $\phi_i(\cdot; \theta)$ is nondecreasing in θ, then the extremal equilibria are nondecreasing in θ.*

The above result yields comparative statics for a variety of strategic settings but a limited class of networks. It differs from results that apply to more general network structures but rely on specific settings - usually involving games of strategic complements or linear best responses. The above also demonstrates how we can use symmetries to discipline network heterogeneity without completely disposing of it. The quotient matrix makes it immediately obvious what restrictions on network structure will collapse the multi-dimensional best-response equation into a one-dimensional one.

A limitation of the above result is that it applies only to one-factor networks. These networks may occur in economic scenarios where the connection from any group i to j can be written multiplicatively. For example, certain stochastic-block models, or in markets where all bilateral exposures are intermediated by one common pool (e.g., unsecured overnight lending). It also nests symmetric benchmarks such as vertex-transitive graphs. Some networks may also be nearly separable across orbits, so that the quotient matrix is close to rank one. In this case, the scalar fixed point and its comparative statics can be used approximate the full game's. We may thus take the exact rank-one case as a tractable benchmark.

Whilst this result guarantees the existence of pure-strategy orbit-invariant equilibria on one-factor networks, it does not establish that these are the only equilibria.[19] This is similar to Galeotti et al. [2010], an influential work which studies sub-classes of equilibria where all agents of the same degree take the same action. Moreover, whilst other non-orbit-invariant equilibria may exist, they need not possess the same favourable comparative statics.

We apply our results to an example: a best shot public goods game with non-linear best responses and convex costs. Best shot public goods games are classic example of a game of strategic substitutes, and have been widely studied in the network games literature.[20] The literature has largely focused on cases with linear best responses. All calculations for the example below can be found in the appendix. These comparative statics cannot generally be obtained for equilibria which are not orbit-invariant, nor for networks which are not one-factor.

Example 1 (Public goods games with non-linear best responses)**.** *We consider a public goods setting on a network where agent i enjoys the public good if either they or one of their neighbours provides it. The outcome of each node's public good provision is stochastic. .*

[19] Indeed, determining for which regions of the parameter space orbit-invariant equilibria are the only equilibria is an important line of future work for this project. It appears that the equilibria of certain network games are typically orbit-invariant over some part of the parameter space, and then undergo a pitchfork bifurcation into an orbit-invariant and a non-orbit-invariant branch.

[20] Notable examples include Bramoullé and Kranton [2007], Bramoullé et al. [2014], and Galeotti et al. [2010].

Agent i chooses effort level $x_i \geq 0$ and produces the public good with probability

$$\Pr[\textit{Public good} = 1](x_i) = 1 - e^{-x_i},$$

producing effort x_i has a quadratic cost equal to $\frac{c}{2}x_i^2$, where $c > 0$. Agent's total utility is therefore

$$u(x_i; x_{-i}) = V \cdot \left(1 - e^{-(x_i + \sum \Gamma_{ij} x_j)}\right) - \frac{c}{2}x_i^2.$$

The best response function is therefore

$$\phi_i(x_i; x_{-i}) = W_0 \left(\frac{V}{c} \cdot e^{-\sum \Gamma_{ij} x_j}\right),$$

where $W_0(\cdot)$ is the principal branch of the Lambert W function. This is the inverse of the equation $f(y) = ye^y$ for $y \geq 0$. It is a non-linear function that satisfies $W_0'(y) > 0$,so that $\partial \phi_i / \partial x_j < 0$ for any $j \in N_i$ (since $e^{-\sum \Gamma_{ij} x_j}$ is decreasing in x_j).

Best responses are therefore non-linear strategic substitutes. W_0 is continuous and bounded from above by $\ln[(2\frac{V}{c} + 1)/(1 + \ln(\frac{V}{c} + 1))]$, so we may apply Proposition 3 to one-factor networks to obtain the following intuitive comparative equilibrium characterisations:

- *There always exists an orbit-invariant equilibrium, and there is a highest and lowest orbit-invariant equilibrium action.*
- *Increasing V/c increases efforts on each orbit.*
- *Adding links to Γ decreases equilibrium actions.*

Orbit-based comparative statics on network structure. The previous section dealt with finding comparative statics on the parameters of the network game. We now shift our focus to comparative statics on structural changes to the network (whilst holding parameters fixed). Changes to the network structure involve adding or removing new links or nodes. We propose a novel method for comparative statics on structural changes to the network: compare networks with a fixed number of orbits, instead of a fixed number of nodes. Any changes to the network must not change the number of orbits of its symmetry group. This approach allows us to change the network in a way that keeps the number of distinct network positions (and thus the number of strategically distinct agents) the same. We call such changes *orbit-preserving changes*:

Definition 8 (Orbit-Preserving Change)**.** *For a network $\Gamma(E, V)$, we say that the addition of vertices V' and edges $E' \subseteq (V' \cup V) \times (V' \cup V)$ to form $\Gamma'(V \cup V', E \cup E')$ is an* orbit-preserving change *if Γ and Γ' have the same number of orbits.*

We illustrate some orbit-preserving changes to the network in Figure 2 in Figure 5.

Orbit-preserving changes contrast with the standard approach of the current literature, which fixes the network size and attempts comparative statics on arbitrary edge changes.

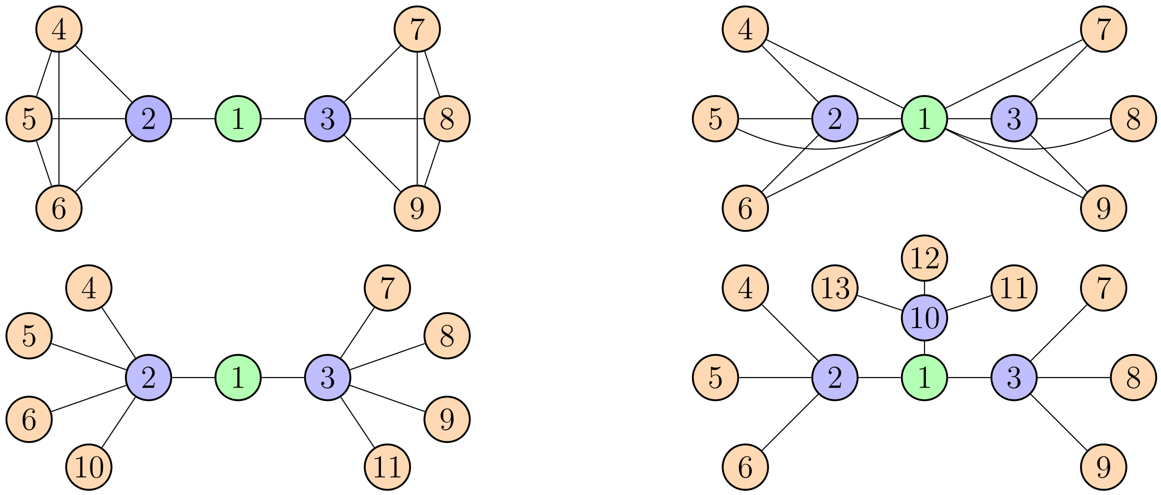

Figure 5: Orbit preserving edge additions (first row) and node additions (second row) to the network depicted in Figure 3.

However, this approach is often inconclusive, even in games of strategic complements, where results regarding supermodular games apply and we may often rank agents' equilibrium actions using centrality measures. This is because the lattice of networks,[21] does not admit strict comparisons between networks with the same number of edges, nor different numbers of nodes.

A limitation of this approach is that edge and node additions must preserve orbit structure. This restriction is not unreasonable if a symmetric network-formation process is assumed. Moreover, the analysis applies in settings where the equilibrium is orbit-invariant —for example, when it is unique or when payoffs exhibit strategic complements. While supermodular-games methods provide some comparative-statics results (since networks form a lattice), our approach remains useful in cases those methods do not directly address, such as comparisons across networks with different numbers of nodes or with the same numbers of nodes and edges but different structures.

We illustrate the advantage of this approach with an example of Hotelling price competition on a network. One-orbit (i.e., vertex-transitive) networks behave like the standard circular Salop model with quadratic costs; indeed, the Salop model is a special case. Below, we show that, even for two-orbit networks, we can find interesting comparative statics on orbit-preserving edge and node additions.

Example 2 (Two Orbit Hotelling)**.** *We consider firms indexed by i embedded on a network $\Gamma = (E, V)$. Each undirected edge connecting two firms represents a mass of 1 consumers who are uniformly distributed on the interval $[0, 1/E]$. Firm i chooses a price p_i which they will charge in all markets (i.e. they cannot choose market-specific prices). Costs are identical for each firm and quadratic: $c(q) = kq^2/2$, where $k > 0$.*[22]

[21] The lattice formed on the set of networks with n nodes ordered by number of edges.

[22] Here we deviate from the standard Hotelling model wherein costs are usually linear. However, strictly

Firms i and j compete for market share on edge $\{i, j\}$ in the standard Hotelling manner. A consumer located at position x on edge i, j has utility $u = v - p_i - tx$ if they purchase from firm i, where $t > 0$ is the constant transport cost faced by consumers and $v > 0$ is the valuation of consumption. Thus, on each edge incident to it, firm i faces linear demand $\frac{1}{2t}(p_j - p_i + t/|E|)$ and so total demand for firm is given by

$$q_i(\mathbf{p}) = \sum \Gamma_{ij} \left(\frac{p_j - p_i + t/|E|}{2t} \right),$$

where $\mathbf{p} \in (\mathbb{R}^+)^n$ is the vector prices charged by firms.

For convex costs, there exists a unique equilibrium price $\mathbf{p}$. Because the equilibrium is unique, by Proposition 1, we have that the equilibrium price vector must be constant on each orbit of Γ so for any orbit σ $p_i^* = p_j^*$ $\forall$ $i, j \in \sigma$. So we may employ our toolkit of orbit-based comparative statics to yield some novel insights. In particular, we focus on the two-orbit setting, which is essentially equivalent to a two-player game on a weighted and directed network. We may therefore reduce out problem an implicit two-equation system which can be easily differentiated to yield the following:

- Adding nodes increases competition: *if we add nodes in a way that preserves the two orbit structures then $\frac{\partial p_i^*}{\partial|\sigma^i|}$ and $\frac{\partial p_i^*}{\partial|\sigma^j|}$ are both < 0.*[23]

- Cross orbit densification increases competitive pressure more than intra-orbit densification: *suppose we add links to the network in a way that preserves the two orbit structures: then $|\frac{\partial p_i^*}{\partial q_{ij}}| > |\frac{\partial p_i^*}{\partial q_{ii}}|$ and $|\frac{\partial p_i^*}{\partial q_{ij}}| > |\frac{\partial p_i^*}{\partial q_{jj}}|$.*

Although we can solve for equilibrium prices using linear algebra, this explicit solution it does not readily yield comparative statics on edge and node additions. The equilibrium-price equation is opaque and highly dependent on the specific network structure—because it involves inverting the adjacency matrix—so the resulting comparative statics are correspondingly network specific. We can also compare two networks of the same size but differing structures, which is usually difficult to achieve with standard methods. This is because these approaches rely on the set of networks on n nodes forming a lattice, but not all networks are strictly ordered.

4 Network Targeting Interventions

Understanding the optimal policy for welfare-maximising interventions in network games is very important. A key area of interest in economic network theory is leveraging the network's structure to maximise the impact of welfare maximising interventions. We build on a recent and influential work, Galeotti et al. [2020], which shows that the optimal intervention for

convex costs couple firms production decisions across markets and thus allow network effects to bite. The insights of this model go through for any convex cost functions.

[23] This assumes that the quotient graph remains unchanged- only orbit sizes increase.

the standard quadratic-linear utility model is pinned down by the network's *spectrum* (i.e. the eigenvalues and associated eigenspaces of the adjacency matrix). We show that the symmetries of the network determine the structure of this spectrum, using a suprising group theoretic result. In this section, we lay out Galeotti et al. [2010]'s model and results, then show how we can use symmetries to gain new structural insights for these results.

4.1 Model

We consider a set of agents indexed $N = \{1, \ldots, n\}$ embedded on an unweighted, undirected and connected network Γ.[24] As a reminder, we abuse notation and also refer to the adjacency matrix of the network as Γ, with entries Γ_{ij}. The adjacency matrix Γ is a symmetric matrix because the network is undirected.

Each agent i simultaneously chooses an action $a_i \in \mathbb{R}$ to maximise their utility function U_i:

$$U_i\left(a, \Gamma\right) = a_i \left(b_i + \beta \sum_{j \in N} \Gamma_{ij} a_j\right) - \frac{1}{2} a_i^2. \tag{1}$$

Here $b_i \in \mathbb{R}$ is i's independent marginal return to their action which does not depend on the action of others. We let $b \in \mathbb{R}^n$ denote the vector of each agent's independent marginal returns. Conversely, $\beta \in \mathbb{R}$ governs returns to the actions of others. Note that if $\beta > 0$, the game exhibits strategic complements, and if $\beta < 0$, it exhibits strategic substitutes.

The Nash equilibrium profile of the game $a^* \in \mathbb{R}^n$ satisfies:

$$a^* = \left[\mathbf{I} - \beta \Gamma\right]^{-1} b,$$

where the inverse in the above is well defined as long as we make the assumption that the largest eigenvalue of Γ is less than $1/|\beta|$.

The utilitarian social welfare at equilibrium is defined as the sum of the equilibrium utilities:

$$W(b, \Gamma) = \sum_{i \in N} U_i\left(a^*, \Gamma\right) = \sum_{i \in N} \frac{1}{2} \left(a_i^*\right)^2 = \frac{1}{2} \left(a^*\right)^T a^*, \tag{2}$$

where the second equality is also a standard result.

Our social planner's problem closely follows that of Galeotti et al. [2020],[25] which characterises the optimal intervention of a planner who acts by changing the *status quo vector* of marginal returns $\hat{b}$ to a vector b, subject to a budget constraint C on the cost of their actions.

The intervention occurs before agents simultaneously choose their actions. The planner's

[24] The unweightedness assumption can be dropped, but makes our exposition clearer.

[25] We have omitted Galeotti et al. [2020]'s non-strategic, non-network spillovers.

targeting problem is given by:

$$\max_b W(b,\Gamma)$$
$$\text{s.t. } a^* = [\mathbf{I} - \beta\Gamma]^{-1} b \text{ and } \sum_{i \in N} \left(b_i - \hat{b}_i\right)^2 \leq C.$$

An important distinction between our work and Galeotti et al. [2020] is that we restrict our attention to unweighted networks (i.e. the entries of the adjacency matrix Γ are either 0 or 1) whilst Galeotti et al. [2020] allows for arbitrary weights in $\mathbb{R}^+$. However, we would not expect an arbitrary network weighted in $\mathbb{R}^+$ to exhibit non-trivial symmetries, as it is unlikely that any two links would have the exactly same weight.[26] Whilst our results in this section would hold for positive weighted networks if we extended Definition 1 to require that a symmetry preserves adjacency and edge-weights,[27] we restrict our attention to unweighted networks where symmetries are most salient. These occur in important economic scenarios such as friendship and collaboration networks, where datasets are usually binary, as well as product relatedness networks.

Eigenvalues and eigenspaces of undirected networks. Galeotti et al. [2020] show that the social planner's problem is easily solved in terms of the matrix Γ's eigenvalues and corresponding eigenspaces. Because we assume that the network is undirected, its adjacency matrix is a real symmetric matrix and therefore diagonalisable. Recall that a matrix Γ is diagonalisable if there exists an orthogonal matrix U consisting of eigenvectors of Γ such that $U\Lambda U^T = \Gamma$, where Λ is a diagonal real $n \times n$ matrix consisting of the eigenvalues of Γ.

We order the eigenvalues of Γ in weakly descending order so that $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$. Note that there may be some repeated values in this chain. We order an orthonormal basis consisting of eigenvectors $u^1, \ldots, u^n$ of Γ according to the magnitude of their associated eigenvalue.[28] If an eigenvalue repeats then its algebraic multiplicity equals geometric multiplicity.[29]

Moreover, let $n - m + 1$ be the dimension of the eigenspace associated with the smallest eigenvalue, λ_m, so that the eigenbasis for this space is $\{u^m, \ldots, u^n\}$. Note that if $m < n$, then $\lambda_m = \lambda_{m+1} = \ldots = \lambda_n$. For ease of reference we shall refer to the smallest eigenvalue as λ_m.

[26] Indeed, important future work is to relax the notion of symmetry to extend our results to "almost symmetric" networks, although symmetries in the underlying network formation process would result in exact symmetric weighted networks.

[27] Formally we would say that if Γ is a weighted graph with weight function $w : E(\Gamma) \to \mathbb{R}^+$, then Π is a symmetry if $\{i,j\} \in E \iff \{\Pi(i), \Pi(j)\} \in E$ and $w(\{i,j\}) = w(\{\Pi(i), \Pi(j)\})$.

[28] Note that when an eigenvalue has geometric multiplicity > 1, the choice of basis orthonormal basis is not unique, but the eigenspace in $\mathbb{R}^n$ clearly is.

[29] If an eigenvalue has multiplicity μ so that $\lambda_i = \lambda i + 1 = \lambda_{i+\mu}$, then the μ corresponding eigenvectors may be ordered randomly without loss of generality.

4.2 Symmetry Structures in Targeting Interventions

4.2.1 Galeotti et al. [2020]: key results and their implications

Galeotti et al. [2020]'s seminal work demonstrates how the eigenvalues and eigenvectors of the matrix Γ cleanly characterises the planner's intervention. Essentially, each of the eigenvectors of Γ describes an independent *principal direction* in which the planner would like to intervene: a direction of intervention which shifts equilibrium actions in proportion to the entries of that particular eigenvector.[30] Their characterisation supplies a transparent spectral formulation of the planner's targeting problem in network games. It has therefore had a wide-reaching economic impact because it yields implementable guidance by solving a typically messy problem very cleanly.

Galeotti et al. [2020] assume that all n eigenvalues of Γ are unique, i.e. *simple*.[31] However, simple eigenvalues turn out to have a surprising implication for a network's symmetry structures. Galeotti et al. [2020] study arbitrary weighted networks that will generally have no non-trivial symmetries. However, this is not the case our setting of unweighted networks with non-trivial symmetries. In particular, assuming simple eigenvalues limits the kinds of symmetry structures which we may study, owing to the following theorem:

Theorem 1. *([Biggs, 1993, Ch.15, Theorem 4]) If all the eigenvalues of a network* Γ *are simple, then every symmetry* Π *of* Γ *satisfies* $\Pi^2 = \mathbf{I}$.

The result above already demonstrates how symmetry groups interplay with network properties - understanding the ramifications for a simple eigenvalues assumption would be impossible without referring to the symmetry group. The seemingly benign assumption that all eigenvalues are simple actually implies that the network must display a specific kind of symmetry: only the set of reflections. This is limiting for two reasons: firstly, many real-world graphs do display large and arbitrary symmetry groups. Secondly, this rules out classes of networks which are economically relevant: for example, those which contain large cliques or demonstrate circular symmetry. Indeed, our running example has the circular symmetry $4 \mapsto 5 \mapsto 6 \mapsto 4$, which does not satisfy this requirement.

Therefore, before we progress with our analysis of the structure of the social planner's results, we must extend the description of the optimal intervention Galeotti et al. [2020] to settings where eigenvalues are not simple.

Let the intervention undertaken by the social planner be denoted by $y^* = b^* - \hat{b}$. Moreover, let $\underline{v} = U^T v$ denote change of basis of the vector $v \in \mathbb{R}^n$ from the standard basis to the eigenbasis. The l^{th} entry of $\underline{v}$ is clearly $v \cdot u^l$. We then have that that the optimal intervention can be determined quite neatly in terms of the spectrum of Γ:

[30] We would not expect this to be the case for an arbitrary change: i.e. if we apply the matrix Γ to a change y we do not expect the resulting vector to be a linear rescaling of y. This is a special property of vectors which lie entirely within a particular eigenspace.

[31] This forms second part of what is called Assumption 2 in Galeotti et al. [2020]. An eigenvalue is called simple if the dimension of its associated eigenspace (the geometric multiplicity) is 1.

Lemma 2 (Extension of Galeotti et al. [2020]). *Let the matrix Γ have eigenvalues $\lambda_1, \ldots, \lambda_n$, ordered in descending order, with associated eigenvectors $u^1, \ldots, u^n$. Let the status-quo vector of agents stand-alone returns be given by $\hat{b}$. The optimal intervention, y^*, and resulting change in equilibrium actions, Δa^*, is given by:*

$$y^* = \sum_{l=1}^{n} \left(\frac{\hat{b} \cdot u^l}{\mu(1-\beta\lambda_l)^2 - 1} \right) u^l, \ \ \text{and} \ \Delta a^* = \sum_{l=1}^{n} \left(\frac{\hat{b} \cdot u^l}{\mu(1-\beta\lambda_l)^3 - 1 + \beta\lambda_l} \right) u^l,$$

where μ is the Lagrange multiplier in the planner's problem, and $\cdot$ is the standard vector dot-product.

Note that the above is essentially identical to Theorem 1 of Galeotti et al. [2020]. However, the proof is somewhat different as the assumption on the eigenvalues is different. We include the proof in the Appendix for robustness. The intuition for the proof of Theorem 1 of Galeotti et al. [2020] (and our version of it) relies on diagonalising both equilibrium actions and welfare in the eigenbasis of Γ. Thus, any intervention decomposes into orthogonal principal components (eigenvectors), each propagating independently with a simple scalar multiplier - the number in brackets in the two expressions in Lemma 2. Theorem 1 (and our version of it) then characterises the optimal intervention by its similarity to these components. We can see which principal directions u^l constitute the largest portion of y^*: those with more extreme eigenvalues (larger for complements, smaller for substitutes) and those that are more parallel to the initial vector of marginal benefits $\hat{b}$ (i.e. those having a larger dot product with it). This should also make clear why we call the eigenvectors the "principal directions" of y^*: they break down the intervention into independent parts, each of which translates to a linear shift in equilibrium actions.

4.2.2 Amplifying and Redistributive Targeting Interventions

We now introduce our first key result, which shows how the symmetries of the network partition the optimal intervention into two economically interesting components:

Theorem 2. *Suppose Γ has r orbits. The principal directions of the optimal intervention y^* consist of:*

- r amplifying modes *which are* constant *on each orbit σ and linearly increase/decrease the action taken by each agent in σ by the* same quantity*; and*
- $n-r$ redistributive modes *which sum* zero *on each orbit σ and* preserve *the total action by reallocating amongst nodes in the same orbit.*

Moreover, the intervention corresponding to the largest eigenvector, u^1, is always constant on each orbit.

Note that when we say that u^l is constant on each orbit σ we mean that $u^l_i = u^l_j$ for every i, j in σ. Conversely, if we say that u^l sums to zero on each orbit, we mean that $\sum_{i\in\sigma} u^l_i = 0$. Note that by the Perron-Frobenius Theorem, u^1 is always simple for connected networks.

The result above shows that the symmetries of Γ play a crucial role in helping us understand the structure of y^*. Whilst previous work could solve for y^* in terms of the spectrum of Γ, the structure of that spectrum is opaque. Viewing the network through the lens of its symmetries removes this opacity and reveals clear patterns in the structure of the optimal intervention. In particular, they predict when agents in the same network position will be treated similarly and when they will be treated differently.

In loose terms, the network spectrum tells us about how a change to a set of nodes' actions will propagate through the network. Because symmetries pin down a lot of the structure of the network, they also pin down the spectrum structure. The eigenvectors that are fixed by every symmetry of the network represent directions of propagation that cannot "cancel out" across symmetric blocks —rather, they evolve by simple scaling by their eigenvalue. We therefore call these *amplification modes*. Conversely, the eigenvectors which sum to zero represent directions of propagation which change sign or redistribute actions within each symmetric block, and their contributions tend to partially cancel. We therefore call these *redistribution modes*.

We can use the quotient graph Γ/G to determine which eigenvectors correspond to constant modes: they are "lifted" eigenvectors of Γ/G.[32] The word "lift" here means that if λ is an eigenvalue of Γ/G with associated eigenvalue u, and we define $\tilde{u}$ by $\tilde{u}_i = u_k$ if $i \in \sigma^k$, then $\tilde{u}$ is an eigenvector of Γ associated with λ.

The intuition for this result is that it is a de-meaning of the optimal intervention according to the sets of symmetrically equivalent agents.[33] This decomposition is helpful because it also partitions the eigenbasis of Γ into a de-meaned part and a residual part (which a standard de-meaning would not do), and thus allows us to leverage the symmetries of the network to understand the propagation of changes to agents' actions. The proof of this relies on a series of fairly deep results in the algebraic graph theory literature.[34] The most notable way that this decomposition respects the propagation structures of the network is that the eigenspace corresponding to the largest eigenvector, which governs the direction of maximal propagation, must always be entirely constant on orbits and hence lives in this symmetric subspace. We show this in the proof of Proposition 4.

Figure 6 illustrates the decomposition of an optimal intervention of the star graph on four nodes into its four principal directions.[35] The size of the nodes indicates their relative magnitude. Positive interventions are coloured green, negative are red, and interventions with a zero magnitude are indicated with a dashed outline. The principal directions associated with the first and last eigenvalues are both constant on each orbit, the middle two are redistributive, and sum to zero on each orbit. Although we have not marked them here, changes in equilibrium actions due to each intervention directions would simply be linear scalings of y^*, and therefore their relative magnitudes and signs would be the same as those

[32] This is because every eigenvalue of Γ/G is an eigenvalue of Γ.

[33] This special demeaning is sometimes used in fluid mechanics and is called the Reynolds Operation. Note that whilst demeaning y^* along any partition of the nodes is possible, only demeaning along the orbits of the symmetry group respects the eigenbasis and thus gives us information about flows through the network.

[34] See Theorem 3 of Ch 9.3 of Godsil and Royle [2001].

[35] For clarity, we have presented un-normalised principal directions here.

of y^*.

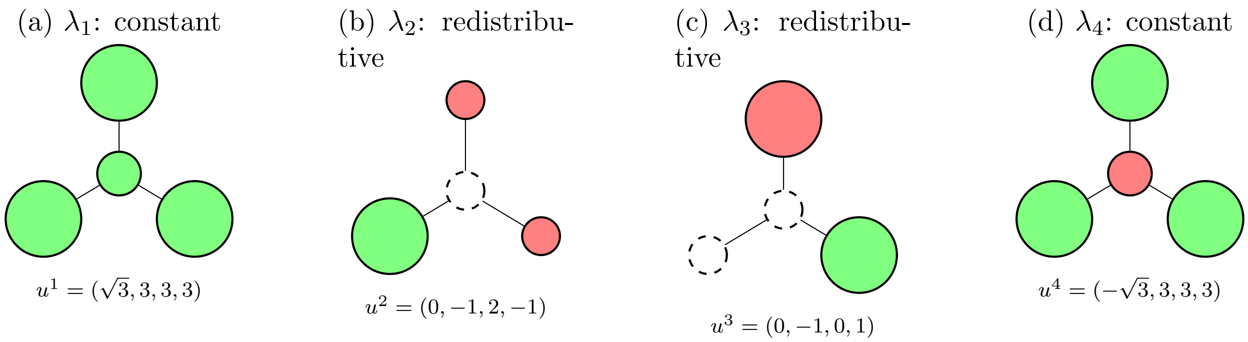

Figure 6: Decomposition y^* into its four principal directions, with associated eigenvectors listed beneath. The total intervention is the sum of these four principal directions.

In some ways, comparing changes in equilibrium action and welfare between nodes in the same orbit is the most natural way to measure inequality whilst keeping network effects fixed. This is because the network's connection structures "looks the same" to nodes in the same orbit, and so the strategic incentives of these nodes (at least in terms of network effects) is the same. This is opposed to nodes in differing orbits, where differences in actions and welfare may simply be due to their occupying very different positions in the network structure.

More symmetric networks with a smaller number of orbits will, paradoxically, by the first point of results of Theorem 2 have fewer eigenvectors in their spectrum which are constant on each orbit. Thus, generally, the social planner's intervention will generally act to redistribute action between agents within each orbit, rather than decreasing/increasing it uniformly across each orbit. Conversely, networks with a larger number of orbits will see interventions which decrease variation within orbits.

The impact of the symmetry structure of a network on variation within orbits is demonstrated in Figure 7 below, where we compare the changes in actions on two very similar networks. Both have the same number of nodes and links, and also the same set of eigenvalues. However, the first network is more symmetric (it has only two orbits and thus two amplification modes) and therefore has more zero-sum eigenvectors in its spectrum (it has six). Conversely, the second network is less symmetric (it has four orbits) and therefore has more constant eigenvectors in its spectrum (it has four). This difference in symmetry structures has a substantial impact on how the optimal intervention acts to create disparities between agents in the same orbit.[36] In the first network, the changes in equilibrium actions within each orbit are more varied, whereas in the second network they are all identical. There are also welfare implications: paradoxically, there will generally be more inequality

[36] The optimal intervention in this example is calculated with identical initial $\underline{b}$ values for both networks, and the magnitude of the projection of the initial $\underline{b}$ onto each eigenspace is identical for each space. $\beta = 0.08$ and $C = 1.5$.

between symmetrically equivalent agents in graphs which are more symmetric that in those which are less symmetric.

Each of the constant and redistribution modes also has an interesting economic meaning as there is interesting interplay between the symmetry of Γ and the types of interventions favoured by the social planner. Firstly, observe that the sum of the redistribution modes represents a balanced budget change to $\hat{b}$ as its entries sum to zero across each orbit (and so in total), and so the resulting total change in equilibrium actions also sums to zero. Thus, this portion of the intervention could be implemented by appropriate taxes and subsidies. Conversely, the constant modes do not would not generically sum to zero and so represent a "pure spending" intervention. Indeed, as Γ becomes increasingly symmetric, in the sense that the number of orbits of G on Γ decreases, the constant modes agree on increasingly many agents and certainly cannot be funded through a balanced budget intervention.

Finally, we note that if Γ has at least one non-trivial symmetry then Γ/G must have dimension $< n$, and so there must be at least one eigenvector which is not lifted from the quotient graph (and zero can never be an eigenvector). This insight leads to the following remark:

Remark 2. *If a node i has no symmetric partners (ie. the orbit containing i consists of only i), then i receives zero intervention under all constant modes.*

So in our running example from the definitions section above, we know that $y_1^0 = 0$, and also $y_2^0 = -y_3^0$. Moreover, this can be seen for the central node in the λ_2 and λ_3 interventions in Figure 6.

4.2.3 Large Budget Interventions

Theorem 2 is especially useful in characterising y^* for games of strategic complements (i.e. $\beta > 0$) in large budget settings. In particular, we find that the social planner incentivises every agent in the same orbit to take increasingly similar equilibrium actions.

To perform this analysis, we extend a proposition from Galeotti et al. [2020] which shows that under large budgets y^* converges to a single eigenspace, to graphs which have arbitrary symmetry groups and thus may have eigenspaces of dimension greater than 1. The statement of the extension and its proof can be found in Section A.3.

Proposition 4. *Let $\beta > 0$. Suppose that i, j are in the same orbit of the symmetry group of the graph. Then as $C \to \infty$, $\frac{y_i^*}{y_j^*} \to 1$ and $\frac{a_i^*}{a_j^*} \to 1$.*

In other words, agents in the same orbit of the graph receive increasingly similar interventions as the budget grows, which leads to them having increasingly similar actions in equilibrium. Note that two agents who are symmetric in their network position may not be symmetric with respect to the game: their marginal returns to their own action, the $\hat{b}_i$ coefficients, may be different. The strongest version of this result would be on a vertex-transitive graph, where every agent would receive an identical intervention, regardless of their $\hat{b}_i$ co-efficients.

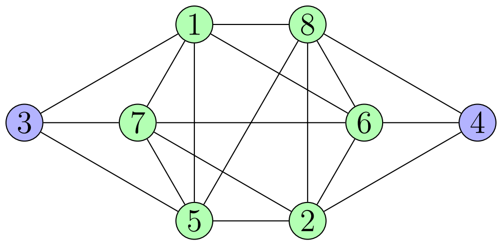

Orbit	i	Δa_i^*
σ^1	1	0.7
	2	0.9
	5	0.7
	6	0.8
	7	0.7
	8	0.8
σ^2	3	0.1
	4	0.2

(a) The optimal intervention in more symmetric networks has more redistributive modes, meaning changes in actions are more varied within each orbit. This network has 2 constant modes and 6 redistributive modes.

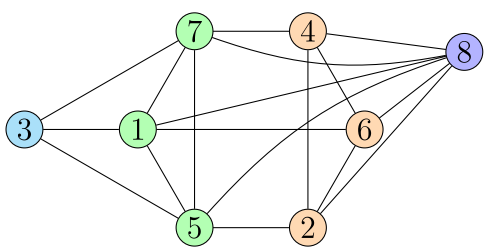

Orbit	i	Δa_i^*
σ^1	1	0.8
	5	0.8
	7	0.8
σ^2	2	0.7
	4	0.7
	6	0.7
σ^3	3	0.1
σ^4	8	0.4

(b) The optimal intervention in less symmetric networks has more amplification modes, meaning changes in actions are more similar within each orbit. This network has 4 constant modes and 4 redistributive modes.

Figure 7: Comparison of identical interventions in two networks with 9 nodes, 18 links and identical eigenvalues but differing symmetry structures. Nodes in the same orbit are the same colour.

The proof of this result combines the characterisation of Theorem 2 with our extension of a result in Galeotti et al. [2020], which shows that under large budgets, the social planner allocates their entire budget to the eigenspaces of the largest or smallest eigenvalues, depending on whether the game exhibits substitutes or complements. Similar to the section above, we extend the result of Galeotti et al. [2020] beyond the simple eigenvalue setting to networks with arbitrary symmetry structures. Our extension can be viewed in the Appendix.

The implications of this proposition are somewhat surprising as it states that, in the presence of increasingly large budgets, the social planner will ignore any differences in $\hat{b}$ and instead focus purely on agents' network position in the underlying graphs This makes somewhat intuitive sense in that the direction of the largest eigenvector maximises flow through the network. However, that the symmetries of the graph are the symmetries of the largest eigenvector may be considered somewhat surprising. This follows from two results in the algebraic graph theory literature.

Firstly, any eigenvector u associated with a simple eigenvalue of Γ must have a particular structure on entries for each orbit of G on Γ: all entries in the same orbit must have the same modulus. This is because if $\Pi \in G$, then if u is a valid eigenvector for λ then so is Πu.[37]. If λ is simple then Πu must be linearly dependent on u, and since λ is real the result follows. If the entries differ in sign then u is a redistributive mode and the entries of each orbit must sum to zero. Thus, the size of the orbit must be even: half its nodes must have positive entry in u and the other half must be negative. So for any two entries i, j in u which are in the same orbit must have the same modulus, but possibly opposite signs. Secondly, the Perron-Frobenius Theorem states that in a strongly connected graph, λ_1 is simple, and all entries for any eigenvector associated with it must have the same sign.

Things, unfortunately, become more complicated when we consider games of strategic substitutes. This is because the smallest eigenvalue is not necessarily simple, as the Perron-Frobenius theorem only characterises the largest eigenvalue of Γ. Thus, the result which states that all entries in the associated eigenvectors must have the same modulus on each orbit no longer applies, nor can we use it to conclude that all entries must have the same sign. However, we can make a complete classification of large budget interventions for two types of networks whose spectrum is very structured: fully vertex-transitive graphs and bipartite graphs.

If we assume that the network is completely symmetric (so that all nodes lie in the same orbit), we may nail down the spectrum of the graph more precisely. We call such graphs *vertex-transitive* graphs. As the result below shows in such graphs, the "opposite" behaviour as described in Proposition 4 seems to occur in that the entirety of the intervention now lies in the redistributive mode of intervention:

Proposition 5. *Let $\beta < 0$. Suppose that Γ is a vertex-transitive graph. Then as $C \to \infty$, $\sum_i y_i^* \to 0$ and $\sum_i a_i^* \to 0$.*

In a sense, the optimal intervention is now maximally asymmetric, in that the nodes in each orbit receive zero net interventions. This is a strong contrast with Proposition 4, where

[37]Notice $\Gamma\Pi u = \Pi\Gamma u = \lambda\Pi u$

the intervention is maximally symmetric (every agent receives an identical intervention). When the network is completely symmetric, then when $\beta > 0$, the social planner wants to give all agents an increasingly identical intervention, which cannot be implemented with a balanced budget. However, when $\beta < 0$, the social planner wants to give agents increasingly dissimilar interventions (in the sense that agents in the same orbit may receive interventions of opposite sign), which can be implemented with a balanced budget intervention.

Notice that Proposition 5 combined with Proposition 4 also suggests some interplay between the symmetry of optimal interventions and whether the game has strategic complements or substitutes. The optimal intervention may shift between a constant intervention mode and a redistributive mode, depending on whether the game has strategic complements or substitutes. This suggests that if we allowed for subsidies and taxes on $\hat{b}$, the social planner could always perform better in games of strategic substitutes than strategic complements, as they may implement the solutions to these in a balanced-budget fashion.

The reason this result holds for vertex-transitive graphs is that only the eigenvector associated with the largest eigenvalue, u^1, is constant on every node in the graph (as the quotient graph Γ/G consists of a single node, and the largest eigenvalue must always be an eigenvalue of the quotient graph). The remaining eigenvectors must always sum to zero on each orbit, and so, in particular, any eigenvectors associated with the smallest eigenvalue, λ_m.

This dramatic difference between complements and substitutes is illustrated for the cycle network on 6 nodes (clearly a vertex-transitive network) in Figure 8. The left-hand panel plots the six entries for y^* (normalised by $\sqrt{C}$) for when the game has strategic complements ($\beta = 0.25$) as a function of increasing budget. Here, as per Proposition 4, the intervention for each agent converges to the same value (roughly 0.4). The equal and opposite divergence for the low-budget entries is caused by the fact that, as discussed above, all eigenspaces but the one associated with the largest eigenvalue are redistribution modes, which still feature in lower-budget interventions. Conversely, the right-hand panel plots the six normalised entries for y^*, as well as their sum, for when the game has strategic substitutes ($\beta = -0.25$). Here, half the nodes converge to the same positive value (roughly 0.4) and the other half converge to the negative of that value (roughly -0.4), so that the sum over all nodes converges to zero (as marked by the dashed line). Hence, the change from complements to substitutes leads to dramatically different behaviour.

It may be tempting to conclude that the modes of intervention associated with the smallest eigenspace always sum to zero over each orbit. However, this is not the case. The result below, however, shows that the symmetry structures of the network play an important role in determining whether y^* is a constant mode or a redistribution mode for large budgets when the game has strategic substitutes ($\beta < 0$). We are able to make this characterisation for *bipartite* graphs: networks whose nodes can be partitioned into two disjoint sets, such that no two nodes that are neighbours in the network are in the same set. For connected networks, this partition is unique.

Proposition 6. *Suppose that Γ is a bipartite network with independent sets R and S. Then if $\beta < 0$, $\lim_{C\to\infty} \sum_{i\in\sigma} y_i^* = 0$ on every orbit σ $\iff$ $\exists \Pi$ such that $\Pi(r) = s$ for some $r \in R$*

(a) For $\beta > 0$: y_i^* converges to a single value.

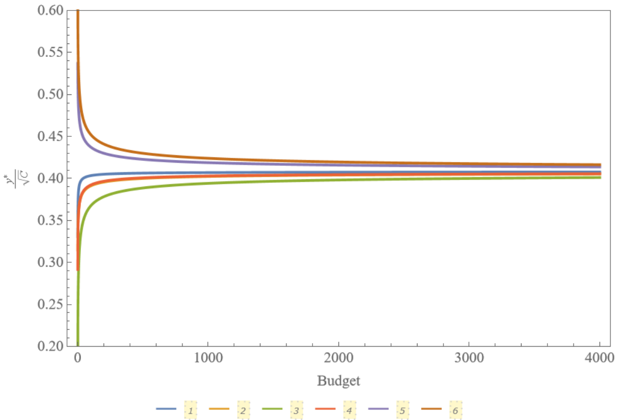

(b) For $\beta < 0$: sum (dashed line) of y_i^* converges to 0.

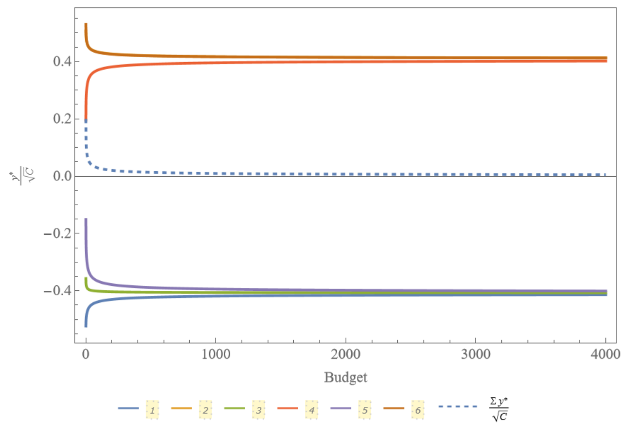

Figure 8: Comparisons of individual y_i^* on a vertex-transitive 6-node network demonstrating convergence to a single value, and a zero-sum for $\beta < 0$.

and $s \in S$. If such a Π exists, it is a bijection from R to S.

This result implies that the social planner should implement redistributive interventions only when there exist two agents in the same orbit (i.e., for whom the network "looks the same") but located in distinct connected components. A bipartite network often represents an economic environment with two types of agents who never interact with agents of the same type: for example, suppliers and buyers. In this case, the result states that it is optimal for the planner to employ a redistributive intervention when the two types are copies of each other in terms of their network incentives. In other words, for every agent of one type, there exists an agent of the other type who occupies an identical position in the network. Note that, because Π is a bijection, the two bipartite sets must be of the same size. Indeed, it follows that Γ/G cannot be bipartite.[38]

For games of strategic substitutes on networks, it is intuitively clear that it is efficient for agents to take different actions from their neighbours (for example, if agent i takes a high action, then their neighbours should take a low action). This is most easily achieved in a bipartite graph, where we may separate the agents into two sets that have no direct links, and thus no direct strategic interactions, with one another. Hence, it is efficient for all nodes in the same bipartite set to take a similar action, as there are no direct spillovers between them. Thus, if two nodes are in different bipartite sets but in the same orbit, the planner wants to increase the differences between them - and hence favours redistributive interventions. This is illustrated in Figure 9 for two bipartite graphs both with 8 vertices and 8 edges, when $\beta = -0.25$.[39] For the first panel, the network is the cycle on eight nodes and is therefore vertex-transitive, so all nodes are in the same orbit. Here, the individual

[38] Although the definition of bipartite graphs for weighted graphs is slightly more involved.

[39] $\hat{b}$ was once again chosen so that each of its projections onto the individual eigenspaces has the same magnitude.

y_i^* values (normalised by $\sqrt{C}$) are plotted as solid lines. It is clear that half of the nodes converge to about 0.38, the others to -0.38, and so the total converges to zero. For the second panel, none of the four orbits contains nodes from two different bipartite sets. Thus, the individual values of y_i^* converge to different values on each orbit: 0.5 for σ^1, 0.1 for σ^2, -0.3 for σ^3, and -0.4 for σ^4.

Although this result currently only applies to bipartite networks, it, combined with Proposition 5, gives us some insight as to what a classification for a wider class of networks might look like. Suppose that the network is partitioned into sets of agents who are not neighbours, and these sets have maximal size.[40]Then the social planner should prefer to give redistributive interventions in large budget scenarios: when there are two agents in independent components of the graph but in the same orbit.

It turns out that the requirements for the smallest eigenvector of a bipartite network to sum to zero on each orbit pin down what kind of symmetries the network can exhibit.

Corollary 2. *If Γ is bipartite and $\lim_{C\to\infty} \sum_{i\in\sigma} y_i^* = 0$ on every orbit σ then $|\sigma|$ must be even for every orbit σ, and G must contain an element Π such that $\Pi^2 = 1$.*

Proof. Firstly, because $\mathbf{u}^n$ has no zero entries, it must be that if $\sum_{i\in\sigma} \mathbf{u}^n = 0$ for every orbit σ, there must be no singleton orbits - which means no node is fixed by every $\Pi \in G$. Moreover, it must be that $|\sigma|$ is even for every orbit, as all entries of $\mathbf{u}^n$ must have the same (non-zero) magnitude on each orbit and must sum to zero. Thus, by the Orbit-Stabiliser Theorem, we must have that 2 divides $|G|$, and so by Cauchy's Theorem, there must be an element of G of order 2. $\square$

The interpretation of this result is that we can essentially pair up all the nodes in one bipartite set with nodes in the other bipartite set. Economically, this means a planner who targets the smallest eigenvector —often the cost-minimising, budget-balanced direction for taxes, subsidies or shocks—can implement it with offsetting transfers inside every class of identical agents and apply the same rule on both sides of the market.

As a final remark, we note that even if we assume that the smallest eigenvalue is simple, we cannot immediately determine whether the optimal result lies in y^c or y^0 (as the sign of all entries need not be the same), but can categorize the convergence of the relative moduli of the components of the optimal intervention:

Remark 3. *Let $\beta < 0$. Moreover, suppose the smallest eigenvalue of Γ, λ_n, is simple. If i, j are in the same orbit of the symmetry group of the graph. Then as $C \to \infty$, $\frac{|y_i^*|}{|y_j^*|} \to 1$.*

The proof follows the same logic as the proof of Proposition 4, except that we may no longer apply the Perron-Frobenius Theorem to assert that all the entries in u^n must have the same sign. Moreover, note that this does not contradict Proposition 5: it would just require that the number of entries of u^n in each orbit would be even and form pairs of equal magnitude and opposite sign.

[40]These are known as maximal independent sets.

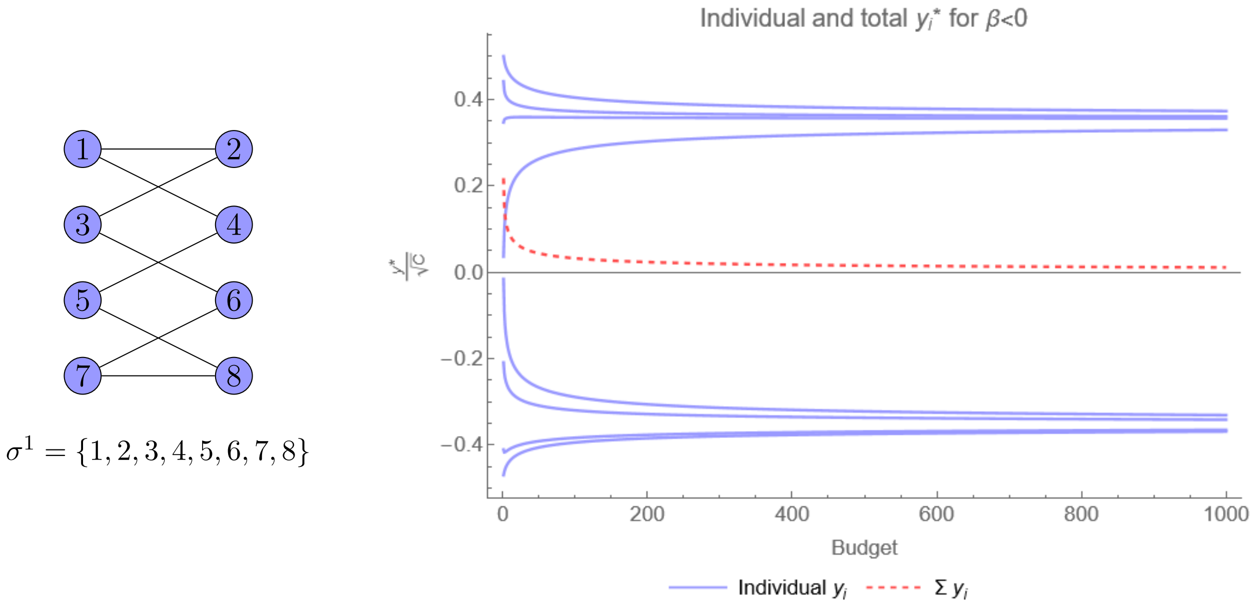

(a) Bipartite network with both bipartite sets in the same orbit: total y_i^* converges to zero.

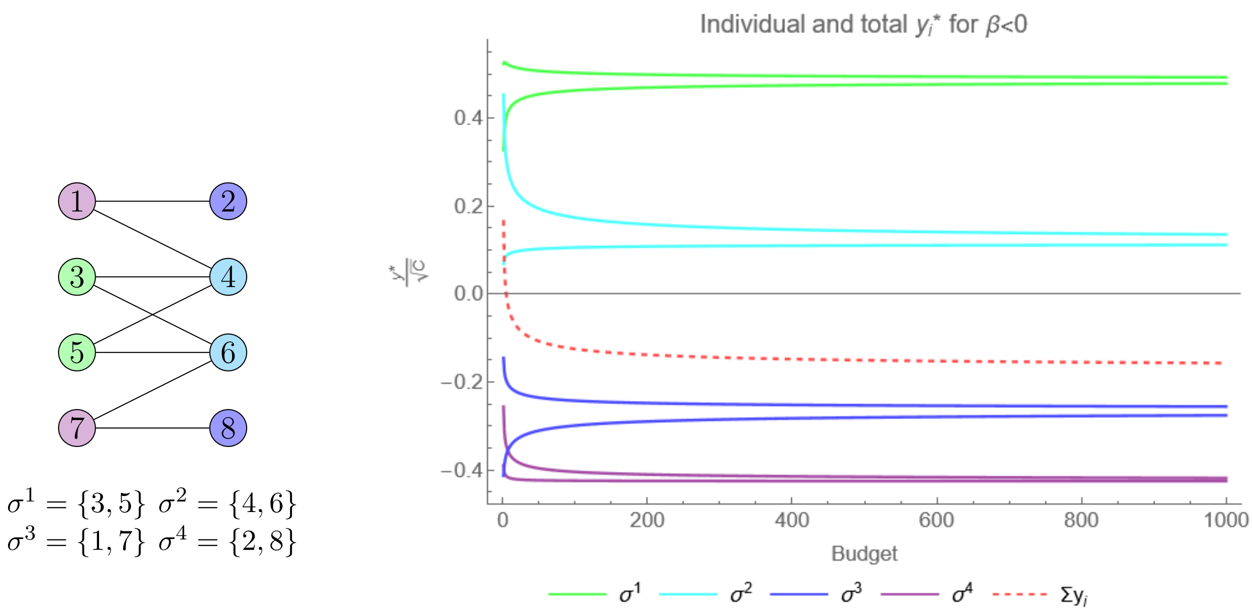

(b) Bipartite network with no orbit containing nodes from different bipartite sets: y_i^* converges to a constant value on each orbit, and does not sum to zero (dotted line).

Figure 9: A comparison of the optimal intervention for strategic substitutes for two bipartite graphs: one's symmetry structure forces a redistributive mode intervention, the other forces a constant mode intervention.

5 Conclusion

This paper illustrates the benefits of a symmetry-based analysis. Our theme is that symmetries provide a useful lens on equilibrium structure by grouping together agents who occupy the same position in the network. We formalise a network's symmetry group and introduce some of its algebraic properties. The first is that it acts on the set of nodes, partitioning them into disjoint, structurally equivalent orbits. The second is that the network's eigenvectors are either constant on orbits or sum to zero on each orbit.

We show that orbits are a network invariant that informs equilibrium structure across a wide variety of games. We prove that many network games admit orbit-invariant equilibria and that, under mild orbit-regularity and neighbour-sum best responses, the highest and lowest such equilibria exist and inherit monotone comparative statics. We demonstrate that symmetries play an important role in optimal targeting for a budget-constrained utilitarian planner, pinning down and refining existing eigenvector-based results. We further show that the number of orbits determines the extent to which an intervention amplifies or dampens actions uniformly across orbits versus redistributes actions across agents. Results are particularly clean for large budgets: interventions become increasingly symmetric for strategic complements but can sum to zero for substitutes. For a budget-constrained utilitarian planner, the optimal intervention decomposes into a uniform component (equal within orbits) and a zero-sum redistribution across orbits: the uniform piece dominates with strategic complements, while redistribution reallocates activity with strategic substitutes. In a Hotelling pricing application, cross-orbit links intensify competition more than intra-orbit links.

The study of symmetry groups of networks is rich and deep yet largely ignored by economists, and this work has only touched on its implications. There are many avenues for future work, some of which we list below.

The first part of this work seeks sufficient conditions for the existence of orbit-invariant equilibria. A valuable exercise would be to determine when games admit only orbit-invariant equilibria and when orbits display heterogeneous actions. This may depend on how strongly the network game amplifies network effects (i.e., how much agents care about what their neighbours do) and on specific geometric aspects of the symmetry group (for example, whether it contains only reflections or also rotations). We can also use the subgroup structure of the symmetry group to classify which symmetries are permitted. For example, a cycle of prime length has no non-trivial symmetric substructures: equilibria must be totally symmetric or totally asymmetric.

A clear next task for the second part of this work is to extend Proposition 5 to a wider class of networks. In particular, it would be interesting to identify when the smallest eigenspace lies entirely in the amplifying mode, as these would correspond to balanced-budget interventions. Another natural application is to study the optimal intervention of a cost-minimising social planner who can levy subsidies and taxes to raise or lower agents' stand-alone returns, rather than a budget-constrained one. There are also results that connect network symmetries to eigenvalues (rather than to the structure of eigenspaces). Since many network results are based on eigenvalues, there may be fruitful applications there as well.

Finally, a natural limitation of this work is that our results have the most bite in networks with non-trivial symmetries. We can justify this as a midway point to discipline network heterogeneity which is not outside the spirit of economic theory. Most economic games assume complete heterogeneity (i.e. there is no network), whilst the current network literature allows arbitrary heterogeneity, in that it does not consider symmetries at all. Symmetries therefore provide a stylisation to improve tractability which is not completely divorced from reality in that real-world networks do appear to display them. However, relaxing our strict symmetry assumptions will serve to make our arguments more robust. There are two avenues through which we might do this. The first is to determine if our result classifying the eigenspace almost holds for totally asymmetric graphs that would, with the addition of a few links, achieve non-trivial symmetries. The second would be to relax the requirement the symmetries are a bijection. This approach can still yield neighbour-preserving mappings on the vertex set which commute with the network's adjacency matrix, and would therefore yield decomposition results for its eigenvectors.

6 Bibliography

References

D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.

M. Akbarpour, S. Malladi, and A. Saberi. Just a few seeds more: The value of network data for diffusion. *American Economic Review*, 115(11):3713–3748, 2025.

N. Allouch. On the private provision of public goods on networks. *Journal of Economic Theory*, 157:527–552, 2015.

N. Allouch and J. Bhattacharya. The key class in networks. *European Economic Review*, 172:104950, 2025.

C. Ballester, A. Calvó-Armengol, and Y. Zenou. Who's who in networks. wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.

A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. The diffusion of microfinance. *Science*, 341(6144):1236498, 2013.

L. A. Beaman. Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the us. *The Review of Economic Studies*, 79(1):128–161, 2012.

N. Biggs. *Algebraic Graph Theory.* Cambridge University Press, 1993.

R. M. Bond, C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489 (7415):295–298, 2012.

Y. Bramoullé and R. Kranton. Public goods in networks. *Journal of Economic theory*, 135 (1):478–494, 2007.

Y. Bramoullé, R. Kranton, and M. D'amours. Strategic interaction and networks. *American Economic Review*, 104(3):898–930, 2014.

A. Calvó-Armengol, E. Patacchini, and Y. Zenou. Peer effects and social networks in education. *The review of economic studies*, 76(4):1239–1267, 2009.

O. Candogan, K. Bimpikis, and A. Ozdaglar. Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905, 2012.

T. Chaney. The network structure of international trade. *American Economic Review*, 104 (11):3600–3634, 2014.

P. K. Chaudhuri, M. O. Jackson, S. Sarangi, and H. Tzavellas. Games under network uncertainty. *arXiv preprint arXiv:2305.03124*, 2024.

P. Dasgupta and E. Maskin. The existence of equilibrium in discontinuous economic games, i: Theory. *The Review of Economic Studies*, 53(1):1–26, 1986.

I. P. Fainmesser and A. Galeotti. Pricing network effects. *The Review of Economic Studies*, 83(1):165–198, 2016.

A. Galeotti, S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv. Network games. *The Review of Economic Studies*, 77(1):218–244, 2010.

A. Galeotti, B. Golub, and S. Goyal. Targeting interventions in networks. *Econometrica*, 88 (6):2445–2471, 2020.

A. Galeotti, B. Golub, S. Goyal, E. Talamàs, and O. Tamuz. Robust market interventions. *arXiv preprint arXiv:2411.03026*, 2024.

C. Godsil and G. F. Royle. *Algebraic Graph Theory*, volume 207. New York: Springer, 2001.

B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.

J. C. Harsanyi and R. Selten. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.

M. O. Jackson and L. Yariv. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98, 2007.

E. Liu and A. Tsyvinski. A dynamic model of input–output networks. *Review of Economic Studies*, 91(6):3608–3644, 2024.

P. Milgrom and J. Roberts. Comparing equilibria. *The American Economic Review*, pages 441–459, 1994.

J. F. Nash. Non-cooperative games. In *The Foundations of Price Theory Vol 4*, pages 329–340. Routledge, 1951.

F. Parise and A. Ozdaglar. Graphon games: A statistical framework for network games and interventions. *Econometrica*, 91(1):191–225, 2023.

A. Plan. Symmetry in n-player games. *Journal of Economic Theory*, 207:105549, 2023.

E. Sadler. Ordinal centrality. *Journal of Political Economy*, 130(4):926–955, 2022.

Y. Zenou and J. Zhou. Sign-equivalent transformations and equilibrium systems: Theory and applications. Working Paper, 2024.

A Appendix

A.1 Additional Definitions

Definition 9 (Group). *A group is a set A together with a binary operation $\cdot : A \times A \to A$ such that the following axioms are satisfied:*

1. ***Closure:*** *For all $a, b \in A$, $a \cdot b \in A$.*
2. ***Associativity:*** *For all $a, b, c \in A$, $(a \cdot b) \cdot c = a \cdot (b \cdot c)$.*
3. ***Identity Element:*** *There exists an element $e \in A$ such that for all $a \in A$, $e \cdot a = a \cdot e = a$.*
4. ***Inverse Element:*** *For each $a \in A$, there exists an element $b \in A$ such that $a \cdot b = b \cdot a = e$, where e is the identity element.*

A.2 Symmetry Structures in Equilibrium - Proofs

A.2.1 Proof of Lemma 1

Proof. Suppose that Π is a symmetry of Γ such that $\Pi(i) = i'$. Consider a strategy profile $\mathbf{x} \in X^n$. Because Π is a symmetry, and there are no self-links in Γ, we have that $N_{i'} = \{\Pi(j) : j \in N_i\}$. So $\{x_{j'}\}_{j' \in N_{i'}} = \{x_{\Pi(j)}\}_{j \in N_i}$, and clearly $|N_i'| = |N_i|$. Because we have assumed that $u_i(x_i, x_{-i}; \Gamma) = u_j(x_j', x_{-j}'; \Gamma)$ if $|N_i| = |N_j|$ and $\{x_k\}_{k \in N_i} = \{x_k'\}_{k \in N_j}$, we clearly have that

$$u_{i'}(x_1, \ldots, x_n; \Gamma) = u_i(x_{\Pi(1)}, \ldots, x_{\Pi(n)}; \Gamma),$$

and so the condition for Π being a symmetry of the game is satisfied. $\square$

A.2.2 Proof of Proposition 1

Proof. This follows from [Plan, 2023, Theorem 2], which states that if Π is a symmetry of a game on a finite player set with strategy space X^n, then if $x^* = (x_1^*, \ldots, x_n^*)$ is an equilibrium, then so is $(x^*_{\Pi(1)}, \ldots, x^*_{\Pi(n)})$. The result then immediately follows from Lemma 1. $\square$

A.2.3 Proof of Remark 1

Proof. If Π is a symmetry of Γ, then we may represent Π with an orthogonal matrix $\mathbf{\Pi}$. We have that Π is a symmetry of Γ iff $\mathbf{\Pi\Gamma} = \mathbf{\Gamma\Pi}$.[41]. But then $\mathbf{\Pi c}^B = \mathbf{\Pi}(\mathbb{I} - \beta\mathbf{\Gamma})^{-1} b\mathbf{\Gamma 1}$. However, clearly $\mathbf{\Pi}^{-1}\mathbf{1} = \mathbf{1}$, and $\mathbf{\Pi}^{-1}\mathbf{\Gamma} = \mathbf{\Gamma\Pi}^{-1}$.[42] Thus, we have that $\mathbf{\Pi c}^B = \mathbf{\Pi}(\mathbb{I} - \beta\mathbf{\Gamma})^{-1}\mathbf{\Pi}^{-1} b\mathbf{\Gamma 1} = \mathbf{c}^B$. $\square$

A.2.4 Proof of Proposition 2

Proof. The sufficient condition comes from Lemma 1. So it remains to prove the necessary condition. Suppose that $S \in \mathcal{G}$ is a symmetry of the game (notice that S need not be a symmetry of Γ). Let $S(i) = i'$ and $S(j) = j'$. We want to show that S is a symmetry of Γ. It is sufficient to show that if $(i, j) \in E(\Gamma) \iff (S(i), S(j)) = (i', j') \in E(\Gamma)$. Consider the strategy profile χ' defined by:

$$\chi_k = \begin{cases} \tilde{x} & \text{if } k = i'; \\ x^1 & \text{if } k \in N_{i'}; \text{ and} \\ x^0 & \text{otherwise.} \end{cases}$$

Denote $\mu(\tilde{x}, x^1) = \mu^1$. Now, because S is a symmetry of the game we must have that for *any* strategy $\mathbf{x} \in X^n$, $u_{S(i)}(x_1, \ldots, x_n; \Gamma) = u_i(x_{S(1)}, \ldots, x_{S(n)}; \Gamma)$. Thus we must have that i''s and i's payoffs must satisfy:

$$\begin{aligned} u_{i'}(\chi'_{i'}, \chi'_{-i'}, \Gamma) &= f\left(\nu(\tilde{x}) + \sum_{j' \in N_{i'}} \mu^1\right) \\ &= u_i(\chi'_{i'}, S(\chi'_{-i'}), \Gamma) \\ &= f\left(\nu(\tilde{x}) + \sum_{j \in N_j} \mu(\tilde{x}, \chi_{S(j)})\right), \end{aligned}$$

where $S(\chi'_{-i'})$ is an abuse of notation, and stands for $(\chi_{S(1)}, \ldots, \chi_{S(n)})_{-i'}$. Because f is an injection, we must have that $\sum_{j' \in N_{i'}} \mu^1 = \sum_{j \in N_j} \mu(\tilde{x}, \chi_{S(j)})$. But $\mu(\tilde{x}, \chi_{S(j)}) \neq 0$ iff $\exists j'$ such that $S(j) = j'$ and $j' \in N_{i'}$. Thus, for the equality to hold, we must have that every neighbour j of i must be the pre-image under S of a neighbour j' of i', and every neighbour j' of i' has a pre-image in N_i. So we have shown that if $(i, j) \in E(\Gamma)$ iff $(S(i), S(j)) \in E(\Gamma)$. $\square$

[41] See [Biggs, 1993, Ch. 15, Prop. 2]

[42] To see this: note that because $\mathbf{\Pi}$ is orthogonal, $\mathbf{\Pi}^{-1} = \mathbf{\Pi}^T$. Thus $\mathbf{\Pi}^{-1}\mathbf{\Gamma} = (\mathbf{\Gamma\Pi})^T = \mathbf{\Gamma\Pi}^{-1}$.

A.2.5 Proof of Proposition 3

Proof. For any symmetric equilibrium, all nodes in orbit σ^r must take the same action x_r. So we may rewrite the BR function ϕ at the orbit-level in terms of the quotient network Γ/G:

$$\Phi_r \left(\sum [\Gamma/G]_{r,s} x_s; \theta_r \right),$$

where from our assumption of agent anonymity we have that $\theta_i = \theta_j = \theta_r \ \forall i, j \in \sigma^r$. Since Γ/G is a rank-one matrix, there exists some $w \in \mathbb{R}^r$ and $\beta \in \mathbb{R}^r$ such that $\Gamma/G = \beta w^T$. Note that Γ/G is non-negative, so w and β are too. Thus, for each orbit, we have that $\sum [\Gamma/G]_{r,s} x_s = \beta_r \sum w_s x_s$. So the role of the network in the BR of any node at some action state x collapses to a single sufficient statistic, $t := \sum w_s x_s$. Substituting $x_s = \Phi_s$ into the expression for t gives

$$t = \sum w_r \Phi_r \left(\beta_r t; \theta_r \right). \tag{3}$$

Since each Φ_r is continuous but for upward jumps, their weighted sum is also continuous in t but for upward jumps. We may therefore apply Corollary 1 of Milgrom and Roberts [1994] to complete the proof. □

A.2.6 Calculations for Example 1

Proof. Solution for best responses The FOC for $u(x_i; x_{-i}) = V \cdot \left(1 - e^{-(x_i + \sum \Gamma_{ij} x_j)}\right) - \frac{c}{2} x_i^2$ with respect to x_i is

$$\begin{aligned} V \cdot e^{-(x_i + \sum \Gamma_{ij} x_j)} - c x_i = 0 \\ x_i e^{x_i} &= \frac{V}{c} \cdot e^{-(x_i - \sum \Gamma_{ij} x_j)} \\ \Rightarrow \phi_i(x_i; x_{-i}) &= W_0 \left(\frac{V}{c} \cdot e^{-\sum \Gamma_{ij} x_j} \right). \end{aligned}$$

Strategic substitutes It is a well-known property of the Lambert-W function that

$$\frac{\mathrm{d}W(y)}{\mathrm{d}y} = \frac{W(y)}{y(1 + W(y))}.$$

Thus

$$\frac{\partial \phi_i(x_i; x_{-i})}{\partial x_j} = -\frac{V}{C} \frac{W_0 \left(\frac{V}{c} \cdot \left(1 - e^{\sum \Gamma_{ij} x_j}\right)\right)}{1 + W_0 \left(\frac{V}{c} \cdot \left(1 - e^{\sum \Gamma_{ij} x_j}\right)\right)} \Gamma_{ij},$$

since $W_0 \geq 0$, we have strategic substitutes.

Upper bound on W_0 If $x = y e^y$ then a well-known bound on W_0 is given by

$$W_0(x) \leq \ln \left(\frac{2x + 1}{1 + \ln(x + 1)} \right),$$

which is increasing in x. For our purposes, $x = \frac{V}{c} \cdot e^{-\sum \Gamma_{ij} x_j}$, which is bounded above by $\frac{V}{c}$ when $x = 0$. Substituting this into the equation above yields the desired upper bound. □

A.2.7 Calculations for Example 2

Uniqueness of equilibrium. We show that the unique equilibrium price is given by

$$\mathbf{p}^* = \frac{t}{|E|}\left(\tilde{\mathbf{D}}^{-1} + \mathbb{I} - \mathbf{D}^{-1}\Gamma\right)^{-1}\mathbf{1}_n, \tag{4}$$

where $\mathbf{D} = \text{Diag}[d_i]$ and $\tilde{\mathbf{D}} = (k/2t)\mathbf{D} + \mathbf{I}$.

Proof. Each firm i solves the problem

$$\max_{pi} p_i \sum_j \Gamma_{ij}\left(\frac{p_j - p_i + t/|E|}{2t}\right) - \frac{k}{2}\left(\sum_j \Gamma_{ij}\left(\frac{p_j - p_i + t/|E|}{2t}\right)\right)^2.$$

This yields the first-order condition

$$\begin{aligned}
&\left(2 + \frac{k}{2t}\sum_j \Gamma_{ij}\right)p_i + \left(\frac{1}{\sum_j \Gamma_{ij}} + \frac{k}{2t}\right)\left(\sum_j \Gamma_{ij}\left(p_j + t/|E|\right)\right) = 0\\
\Rightarrow &\left(2 + \frac{k}{2t}d_i\right)p_i^* = \left(\frac{1}{d_i} + \frac{k}{2t}\right)\sum_j \Gamma_{ij}p_j + \left(1 + \frac{kd_i}{2t}\right)\frac{t}{|E|},
\end{aligned}$$

where d_i is the degree of node i. Writing this as a matrix system yields

$$\begin{aligned}
&\left(2\mathbb{I}_n + \frac{k}{2t}D\right)p = \left(D^{-1} + \frac{k}{2t}\mathbb{I}_n\right)\Gamma p + \left(\mathbb{I}_n + \frac{k}{2t}D\right)\frac{t}{|E|}\mathbf{1}_n\\
\Rightarrow \quad &p^* = \frac{t}{|E|}\left(\tilde{D}^{-1} + \mathbb{I}_n - D^{-1}\Gamma\right)^{-1}\mathbf{1}_n,
\end{aligned}$$

where D is the diagonal matrix consisting of each nodes' degree and $\tilde{D} = \mathbb{I}_n + \frac{k}{2t}D$. Note that the inverse is well defined as $(\tilde{D}^{-1} + \mathbb{I}_n - D^{-1}\Gamma)$ is a symmetric diagonal dominant matrix, as its i^{th} diagonal element is given by $2 + kd_i/2t$, whilst the sum of its off diagonal elements are $1 + kd_i/2t$, which is strictly less. □

Comparative statics calculations:

Proof. Without loss of generality, we focus on comparative statics for the orbit indexed by 1. The following calculations were performed in Wolfram Mathematica.

Let $d_{\sigma^i} = q_{ii} + q_{ij}$ be the degree of σ^i in Γ/G. We have that :

$$p_1^* = \frac{\left(kd_{\sigma^1} + 2t\right)\left(kd_{\sigma^2}\left(d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\right) + 2t\left(d_{\sigma^1}d_{\sigma^2} + d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\right)\right)}{2\left(kd_{\sigma^1}d_{\sigma^2}(q_{12} + q_{21}) + 2t\left(d_{\sigma^1}d_{\sigma^2} + d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\right)\right)\left(d_{\sigma^1}|\sigma^1| + d_{\sigma^2}|\sigma^2|\right)}, \tag{5}$$

which is clearly decreasing in $|\sigma^1|$ and $|\sigma^2|$. □

Proof. Again, we focus on orbit 1 WLOG. We substitute in the feasibility constraint for cross orbit edges $|\sigma^2| = q_{12}|\sigma^1|/q_{21}$ into Equation (5) to obtain:

$$p^*_{\sigma^1} = \frac{q_{21}\big(kd_{\sigma^1} + 2t\big)}{2|\sigma^1|\,\big(d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\big)} \cdot \frac{\big(d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\big)\big(kd_{\sigma^2} + 2t\big) + 2t\,d_{\sigma^1}d_{\sigma^2}}{2t\big(d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\big) + d_{\sigma^1}d_{\sigma^2}\big(k(q_{12} + q_{21}) + 2t\big)},$$

We then have that the difference between increasing links within orbit 1 (increasing q_{11}) and increasing links between orbit 1 and orbit 2 (increasing q_{12} - note that if q_{12} increases then q_{21} must increase if the size of each orbit is to remain fixed.):

$$\frac{\partial p^*_{\sigma^1}}{\partial q_{12}} - \frac{\partial p^*_{\sigma^1}}{\partial q_{11}} = -\,q_{21}\,d_{\sigma^2}\,\big(kd_{\sigma^1} + 2t\big)\cdot \frac{\Big(k^2 d_{\sigma^1}d_{\sigma^2}\big(d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\big)^2 + 2k\,\mathcal{P}_{12}(q;d)\,t + 4\big(d_{\sigma^1}d_{\sigma^2} + d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\big)^2 t^2\Big)}{2\big(d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\big)^2\Big(k\,d_{\sigma^1}\,(q_{12} + q_{21})d_{\sigma^2} + 2\big(d_{\sigma^1}d_{\sigma^2} + d_{\sigma^1}q_{21} + d_{\sigma^2}q_{12}\big)t\Big)^2\,\sigma^1},$$

where we have that:

$$\mathcal{P}_{12}(q;d) = (d_{\sigma^1} + d_{\sigma^2})\big(d^2_{\sigma^1}q^2_{21} + d^2_{\sigma^1}d_{\sigma^2}\,q_{21} + 2d_{\sigma^1}d_{\sigma^2}\,q_{12}q_{21} + d^2_{\sigma^2}q^2_{12}\big) + 2d^2_{\sigma^1}d^2_{\sigma^2}\,q_{12}.$$

Since all terms and factors on the RHS are positive, the difference is negative. Since $\frac{\partial p^*_{\sigma^1}}{\partial q_{12}}$ and $\frac{\partial p^*_{\sigma^1}}{\partial q_{11}}$ are both negative, this implies $\left|\frac{\partial p^*_{\sigma^1}}{\partial q_{12}}\right| > \left|\frac{\partial p^*_{\sigma^1}}{\partial q_{11}}\right|$. □

A.2.8 Proof of Theorem 2

Proof. From Lemma 2, the change in equilibrium action due to the portion of y^* in the direction of u^l is given by

$$[\mathbb{I} - \beta\Gamma]^{-1}\,\underline{y}^*_l u^l = \frac{1}{1-\beta\lambda_l}\left(\hat{b}_l\frac{\alpha_l}{\mu - \alpha_l}\right)u^l,$$

where we define $\alpha_l = 1/(1-\beta\lambda_l)^2$. Thus, if we can show that u^l is either constant on each orbit σ or sums to zero, then the result holds.

Indeed, this follows from a series of results in the algebraic graph theory literature, which hinge on the fact that the orbits $\{\sigma^1, \ldots, \sigma^r\}$ of the symmetry group G form an *equitable partition* of Γ. A partition π of Γ with cells $C_1, \ldots, C_r$ is called equitable if for any cell C_k, we have that for each $i \in C_k$, $|N_i \cap C_l|$ is the same, for all $l \in \{1, \ldots, r\}$. This property allowed us to define the quotient graph in the definitions section. The *characteristic matrix* P of π is the $n \times r$ matrix where $P_{ik} = 1$ if $i \in C_k$.

For any equitable partition, we can define a weighted and directed graph, denoted Γ/π, whose nodes are the cells of π and whose edges are the constant number of links between each cell of π. Notice that for the partition generated by G, this is precisely the quotient graph Γ/G defined in the section above. Let us denote its adjacency matrix by H.

The following pair of lemmas then implies the result:

Lemma 3 (Godsil and Royle [2001], Ch.9, Lemma 3.1)**.** *Let π be an equitable partition of the graph Γ, with characteristic matrix P, and let H be the adjacency matrix of Γ/π. Then $\Gamma P = PH$.*

Lemma 4 (Godsil and Royle [2001], Ch.9, Lemma 3.2)**.** *A partition π of Γ is equitable if and only if the column-space of P is Γ-invariant.*

Recall that the column-space of a $n \times r$ matrix P is invariant with respect to a matrix $n \times n$ Γ if for every column vector p^i of P, Γp^i lies in the span of $p^1, \ldots, p^r$. The proof of Lemma 4 hinges on the fact that the column space of P is Γ-invariant iff there is a matrix H such that $\Gamma P = PH$. But since π is an equitable partition, we may let H be the adjacency matrix of Γ/G. Thus, if u is an eigenvector of H with eigenvalue λ then $\Gamma Pu = PHu = \lambda Pu$, and so Pu is an eigenvector of Γ with eigenvalue λ. Thus, we say that P 'lifts' the eigenvectors of H in $\mathbb{R}^r$ to eigenvectors of Γ in $\mathbb{R}^n$.

If P is Γ-invariant, then so is the matrix formed by the orthogonal complement of the column space of P, $P^{\perp}$. This implies that there also exists a matrix H' such that $\Gamma P^{\perp} = P^{\perp}H'$. The same logic applies to the eigenvectors of H'. Since the dimensions of the column space of H and H' must sum to n, every eigenvector of Γ either has the form Pu, (where v is an eigenvector of H) and is constant on the cells of P, or sums to zero on the cells of P. □

A.2.9 Proof of Proposition 4

Proof. First we show that $\lim_{C\to\infty} \frac{y_i^*}{y_j^*} = \frac{u_i^1}{u_j^1}$. In order to deal with the fact that $y^* \to \infty$ as $C \to \infty$, we consider $\underline{\underline{y^*}} = \frac{y^*}{||y||}$. It is clear from Proposition 7 that as $C \to \infty$, $\rho\left(\underline{\underline{y^*}}, u^1\right) \to 1$. This implies that $\lim_{C\to\infty} \underline{\underline{y^*}} = u^1$, as we have that

$$||\underline{\underline{y^*}} - u^1||^2 = ||\underline{\underline{y^*}}||^2 + ||u^1||^2 - 2\underline{\underline{y^*}} \cdot u^1 = 2(1 - \underline{\underline{y^*}} \cdot u^1)$$

and so for every $\epsilon^2/2$ $\exists \underline{C}$ such that if $C > \underline{C}$, $|\underline{\underline{y^*}} \cdot u^1 - 1| < \epsilon^2/2$, and so $||\underline{\underline{y^*}} - u^1|| < \epsilon$. Hence we also have $\lim_{C\to\infty} \underline{\underline{y_i^*}} = u_i^1$.

We now show that if i, j are in same orbit of G on Γ, then $u_i^1 = u_j^1$. To do so, we need the following lemma:

Lemma 5. *([Biggs, 1993, Ch.15, Lemma 3]) Let λ be a simple eigenvalue of Γ, and let u be a corresponding eigenvector. If the permutation matrix P represents an automorphism of Γ then $Pu = \pm u$.*

Thus if i, j are in the same orbit of Γ, $u_i = \pm u_j$. To see this suppose, that P maps node i to node j then, $u_j = \pm u_i$. But this must be true for all j such that there is a permutation on G that maps i to j: i.e. the orbit containing i. Now, recall that from the Perron-Frobenius Theorem, λ_1 is simple, and all its entries have the same sign. Thus for any automorphism P of Γ, $Pu = u$. Thus, $u_i = u_j$ if i and j are in the same orbit under Γ. □

A.2.10 Proof of Proposition 5

Proof. It is clear that, by following a similar logic to the first part of the proof of Proposition 4, we can show that as $C \to \infty$, $\sum_i y^* \to \sum_i \frac{\sqrt{C}}{||b^{\lambda_m}||} \sum_l \underline{\hat{b}}_l u_i^l$. In particular, for every ϵ/n there exists a $\underline{C}$ such that if $C > \underline{C}$, then

$$
\begin{aligned}
&|y_i - \frac{\sqrt{C}}{||b^{\lambda_m}||} \sum_l \underline{\hat{b}}_l u_i^l| < \frac{\epsilon}{n} \\
&\Rightarrow |\sum_i \left(y_i - \frac{\sqrt{C}}{||b^{\lambda_m}||} \sum_l \underline{\hat{b}}_l u_i^l \right)| < \sum_i |y_i - \frac{\sqrt{C}}{||b^{\lambda_m}||} \sum_l \underline{\hat{b}}_l u_i^l| < \epsilon.
\end{aligned}
$$

It is clear that as $C \to \infty$, $\sum_i y^* \to \sum_i \frac{\sqrt{C}}{||b^{\lambda_m}||} \sum_l \underline{\hat{b}}_l u_i^l$. Therefore, it remains to prove that $\sum_i \left(\sum_{l=m}^n u_i^l \right) = 0$ on a vertex-transitive graph.

Notice that all nodes in a vertex-transitive graph must necessarily have the same degree, as symmetries of the graph preserve the number of neighbours. Let k be the degree of Γ. Furthermore, Γ is vertex-transitive, then Γ/G consists of a single node, with self-degree k. Thus, using the notation from our proof of Theorem 2, the adjacency matrix H of the quotient graph is just the scalar k, and Γ/G has a single eigenvalue of k with associated eigenvector 1.

Now, all regular networks of degree k have largest eigenvalue k with associated eigenvector $k\mathbf{1}$. Since $\Gamma/G = \mathbf{1}$ in a vertex-transitive graph, we have that the only eigenvector of Γ with form Pu is the largest one. Therefore the eigenvectors associated with the remaining eigenvalues $\lambda_2, \ldots, \lambda_m$ must sum to zero on each orbit of G on Γ. Thus, the eigenvectors $u^m, \ldots, u^n$ associated with the smallest eigenvalue must sum to zero. Since a vertex-transitive graph only has a single orbit, we must have that the entries of each of $u^m, \ldots, u^n$ sum to zero. $\square$

A.2.11 Proof of Proposition 6

Proof. In a similar line to our proof of Proposition 5 we want to show that this condition is necessary and sufficient for $\sum_{i \in \sigma} \mathbf{u}_i^n = 0$ for each orbit σ. The proof of this hinges on the spectral structure of bipartite graphs. Suppose Γ has bipartite components of size r and s. Firstly, note that if (u, v) is a length $r + s$ eigenvector of the bi-adjacency matrix of Γ with corresponding eigenvalue λ, then $(u, -v)$ is an eigenvector of $-\lambda$ ([Godsil and Royle, 2001, p.178]). Thus, a bipartite graph has $\lambda_n = -\lambda_1$. Moreover, as a result of the Perron-Frobenius Theorem, λ_n is simple and $\mathbf{u}^n$ has no zero entries.

We now prove the sufficient condition ($\Rightarrow$). If $\exists \Pi$ such that $\Pi(r) = s$ for some $r \in R$ and $s \in S$ then it must be that r and s are in the same orbit σ of G on Γ, which by remark 3 means that $|\mathbf{u}_s^n| = |\mathbf{u}_r^n|$. However, their signs are different: the smallest eigenvector $\mathbf{u}^n$ has the form $(u, -v)$ where the entries of u are positive and correspond to the nodes in R and the entries of v are also positive and correspond to the nodes in S. Thus $\text{sgn}(\mathbf{u}_s^n) = -\text{sgn}(\mathbf{u}_r^n)$. Thus it must be that $\sum_{i \in \sigma} \mathbf{u}_i^n = 0$, which by theorem 2 implies that $\mathbf{u}^n$ sums to zero on every orbit of G.

For the necessary condition ($\Leftarrow$), notice that if $\sum_{i\in\sigma} \mathbf{u}_i^n = 0$ for every orbit σ, then $\mathbf{u}^n$ has entries of opposite sign in every orbit (since none of its entries can be zero), which means that each orbit consists of nodes from both r and s, which implies $\exists \Pi \in G$ such that $\Pi(r) = s$ for some $r \in R$, $s \in S$.

To see the final part of the statement, notice that all of r's neighbours are in S, and all of s's neighbours are in R (by virtue of Γ being bipartite). Hence, if r goes into s's position, all of s's neighbours in R must be mapped to neighbours of r in S. Because Γ is connected, repeating this logic shows that $\Pi(R) = S$. □

A.3 Additional Proofs

The proof of lemma 2 follows from eq. (6) in the proof of proposition 7 below.

We let ρ denote the *cosine similarity* between two non-zero vectors y, z, defined by

$$\rho(y,z) = \frac{y \cdot z}{||y||\,||z||},$$

where $||\cdot||$ is the Euclidean norm on $\mathbb{R}^n$.

We may now proceed with our extension of what is Proposition 1 of Galeotti et al. [2020]:

Proposition 7. *[Extension of Galeotti et al. [2020]] Suppose that Γ is undirected, then*

1. *If $\beta > 0$ then as $C \to \infty$, $\rho\left(y^*, \sqrt{C}u^1\right) \to 1$,*
2. *If $\beta < 0$ then as $C \to \infty$, $\rho\left(y^*, \frac{\sqrt{C}}{||\hat{b}^{\lambda_m}||}\left[\sum_{l=m}^{n} \underline{\hat{b}}_l u^l\right]\right) \to 1$.*

Where $\hat{b}^{\lambda_m}$ is the projection of $\hat{b}$ onto the λ_m-eigenspace.[43] Part 1 of the above remains unchanged from the original statement, as the Perron-Frobenius Theorem states that if a graph is strongly-connected and its adjacency matrix is non-negative then λ_1 is equal to the spectral radius of the graph, and is also simple. We thus only prove part 2, which we exclude from the body text as it follows a similar line of reasoning to that of the original result and therefore makes a limited methodological contribution.

Proof of Proposition 7:

Proof. Showing that $\frac{W^*}{W^S} \to 1$ as $C \to \infty$: Assume $\beta < 0$. Let $\tilde{x}$ be the optimal intervention when the planner is constrained to allocate their entire budget to a vector in the λ_m eigenspace, so that $\tilde{x}_l = 0 \; \forall l \leq m-1$. It is clear that $\tilde{x}_l$ is the same $\forall l \geq m$. This is because, similar to the unconstrained problem, we have that

$$\begin{aligned}
\mathcal{L} &= w \sum_{l=m}^{n} \alpha_l \left(1 + x_l\right)^2 \hat{b}_l + \mu \left(C - \sum_l \underline{\hat{b}}_l^2 x_l^2 \right) \\
\frac{\partial \mathcal{L}}{\partial x_l}\Big|_{x_l^*} &= 2\underline{\hat{b}}_l^2 \left(\alpha_l(1+\tilde{x}_l) - \mu\tilde{x}_l\right) = 0 \\
\Rightarrow \tilde{x}_l &= \frac{w\alpha_l}{\mu - \alpha_l}.
\end{aligned}$$

[43] In other words, it is the length $(n-m+1)$ vector of the form $\sum_{l=m}^{n} (\hat{b}\cdot u^l)u^l$.

But $\alpha_m = \alpha_{m+1} \ldots = \alpha_n$, as all these vectors belong to the same eigenspace. We shall simply refer to α_l for $l \geq m$ as α_m, and similarly for $\tilde{x}_m$ and x_m^*.

Because $\sum \tilde{x}_m^2 \underline{\hat{b}}_l^2 = C$, we have that $\tilde{x}_m = \frac{\sqrt{C}}{||\hat{b}^{\lambda_m}||}$. Notice that $\tilde{x}_m \geq x_m^*$ as

$$\sum_{l=m}^{n} {x_m^*}^2 \underline{\hat{b}}_l^2 = C - \sum_{l=1}^{m-1} {x_l^*}^2 \underline{\hat{b}}_l^2 \quad \Rightarrow x_m^* = \frac{\sqrt{C - \sum_{l=1}^{m-1} {x_l^*}^2 \underline{\hat{b}}_l^2}}{||\hat{b}^{\lambda_m}||} \leq \frac{\sqrt{C}}{||\hat{b}^{\lambda_m}||}. \tag{6}$$

Now, we have that

$$\frac{W^*}{W^s} = \frac{\sum_{l=1}^{m-1} \underline{\hat{b}}_l^2 \alpha_l \left({x_l^*}^2 + 2x_l^*\right) + \sum_{l=m}^{n} \underline{\hat{b}}_l^2 \alpha_m \left({x_m^*}^2 + 2x_m^*\right) + \sum_{l=1}^{n} \underline{\hat{b}}_l^2 \alpha_l}{\sum_{l=m}^{n} \underline{\hat{b}}_l^2 \alpha_m \left(\tilde{x}_m^2 + 2\tilde{x}_m\right) + \sum_{l=1}^{n} \underline{\hat{b}}_l^2 \alpha_l} \tag{7}$$

$$\leq \frac{\sum_{l=1}^{m-1} \underline{\hat{b}}_l^2 \alpha_l \left({x_l^*}^2 + 2x_l^*\right)}{\sum_{l=m}^{n} \underline{\hat{b}}_l^2 \alpha_m \left(\tilde{x}_m^2 + 2\tilde{x}_m\right) + \sum_{l=1}^{n} \underline{\hat{b}}_l^2 \alpha_l} + \frac{\sum_{l=m}^{n} \underline{\hat{b}}_l^2 \alpha_m \left({x_m^*}^2 + 2x_m^*\right) + \sum_{l=1}^{n} \underline{\hat{b}}_l^2 \alpha_l}{\sum_{l=m}^{n} \underline{\hat{b}}_l^2 \alpha_m \left(\tilde{x}_m^2 + 2\tilde{x}_m\right) + \sum_{l=1}^{n} \underline{\hat{b}}_l^2 \alpha_l} \tag{8}$$

$$\leq \frac{\sum_{l=1}^{m-1} \underline{\hat{b}}_l^2 \alpha_l \left({x_l^*}^2 + 2x_l^*\right)}{\sum_{l=m}^{n} \underline{\hat{b}}_l^2 \alpha_m \left(\tilde{x}_m^2 + 2\tilde{x}_m\right)} + 1 \tag{9}$$

$$\leq \frac{\sum_{l=1}^{m-1} \underline{\hat{b}}_l^2 \alpha_l \left({x_l^*}^2 + 2x_l^*\right)}{\sum_{l=m}^{n} \underline{\hat{b}}_l^2 \alpha_m \tilde{x}_m^2} + 1. \tag{10}$$

Equation (9) follows from the fact that $\underline{\hat{b}}_l^2 \alpha_l > 0 \ \forall l$ and $\tilde{x}_m \geq x_m^* \geq 0$. Now, notice that when $\beta < 0$, $\alpha_{m-1} \geq \alpha_l$, and $x_{m-1}^* = \alpha_{m-1}/(\mu - \alpha_{m-1}) \geq \alpha_l/(\mu - \alpha_l) = x_l^*$, for all $l < m-1$. Thus,

$$\begin{aligned}
\sum_{l=1}^{m-1} \underline{\hat{b}}_l^2 \alpha_l \left({x_l^*}^2 + 2x_l^*\right) &\leq \alpha_{m-1} \left({x_{m-1}^*}^2 + 2x_{m-1}^*\right) \sum_{l=1}^{m-1} \underline{\hat{b}}_l^2 \\
&\leq \alpha_{m-1} \left(\left(\frac{\alpha_{m-1}}{\mu - \alpha_{m-1}}\right)^2 + 2\left(\frac{\alpha_{m-1}}{\mu - \alpha_{m-1}}\right)\right) ||b||^2 \\
&\leq \alpha_{m-1} \left(\left(\frac{\alpha_{m-1}}{\alpha_m - \alpha_{m-1}}\right)^2 + 2\left(\frac{\alpha_{m-1}}{\alpha_m - \alpha_{m-1}}\right)\right) ||b||^2 \\
&= \frac{2\alpha_{m-1}(\alpha_m - \alpha_{m-1})}{(\alpha_m - \alpha_{m-1})^2} ||b||^2.
\end{aligned}$$

Substituting this into the W^*/W^s inequality and using the fact that the budget constraint must bind for $\tilde{x}$, we have

$$\frac{W^*}{W^s} \leq \frac{\alpha_{m-1}^2}{(\alpha_m - \alpha_{m-1})^2} \frac{(2\alpha_m - 2\alpha_{m-1})}{\alpha_m C} ||b||^2 + 1.$$

It is clear that as $C \to \infty$, the right-hand side of the inequality converges to 1. Moreover, $1 \leq \frac{W^*}{W^s}$. Thus, applying the Squeeze Theorem yields the desired result.

Showing that $\rho\left(y^*, \frac{\sqrt{C}}{||\hat{b}^{\lambda_m}||}\left[\sum_{l=m}^{n} \underline{\hat{b}}_l u^l\right]\right) \to 1$ *as* $C \to \infty$*:* We have that

$$\begin{aligned}
\rho\left(b^* - \hat{b}, \frac{\sqrt{C}}{||\hat{b}^{\lambda_m}||}\left[\sum_{l=m}^{n} \underline{\hat{b}}_l u^l\right]\right) &= \frac{\left(b^* - \hat{b}\right) \cdot \left(\sum_{l=m}^{n} \underline{\hat{b}}_l u^l\right)}{||b^* - \hat{b}|| \times ||\sum_{l=m}^{n} \underline{\hat{b}}_l u^l||} \\
&= \frac{\left(b^* - \hat{b}\right) \cdot \left(\sum_{l=m}^{n} \underline{\hat{b}}_l u^l\right)}{\sqrt{C} \times ||b^{\lambda_m}||} \\
&= \frac{\sum_{l=m}^{n} \underline{\hat{b}}_l \left[\left(b^* - \hat{b}\right) \cdot u^l\right]}{\sqrt{C} \times ||b^{\lambda_m}||},
\end{aligned}$$

where we have dropped the factor of $\sqrt{C}/||b^{\lambda_m}||$ in the first equality because it cancels in the numerator and the denominator. Note that $b^{\lambda_m} = \sum_{l=m}^{n} \underline{\hat{b}}_l u^l$ in Cartesian co-ordinates.

So it remains to compute $\left(b^* - \hat{b}\right) \cdot u^l$ for each $l \in \{m, \ldots, n\}$. Notice that

$$\left(b^* - \hat{b}\right) \cdot u^l = \left(U\left(\underline{b}^* - \underline{\hat{b}}\right)\right) \cdot u^l = \left(\underline{b}^* - \underline{\hat{b}}\right) \cdot \left(U^T u^l\right) = \left(\underline{b}^* - \underline{\hat{b}}\right)_l = x_l^* \underline{\hat{b}}_l.$$

Moreover, since $x_l^* = x_m^*$ for each $l \in \{m, \ldots, n\}$, we have that

$$\begin{aligned}
x_m^* &= \frac{1}{||b^{\lambda_m}||}\sqrt{C - \sum_{l=1}^{m-1} x_l^{*2} \underline{\hat{b}}_l^{\,2}} \\
&\geq \frac{1}{||b^{\lambda_m}||}\sqrt{C - \left(x_{m-1}^*\right)^2 \sum_{l=1}^{n} \underline{\hat{b}}_l^{\,2}} \\
&\geq \frac{1}{||b^{\lambda_m}||}\sqrt{C - \left(\frac{\alpha_{m-1}}{\alpha_m - \alpha_{m-1}}\right)^2 ||\hat{b}||^2}.
\end{aligned}$$

Plugging this back into the expression for cosine similarity yields

$$\begin{aligned}
\rho\left(b^* - \hat{b}, \frac{\sqrt{C}}{||\hat{b}^{\lambda_m}||}\left[\sum_{l=m}^{n} \underline{\hat{b}}_l u^l\right]\right) &\geq \frac{\sum_{l=m}^{n} \underline{\hat{b}}_l^2}{||\hat{b}^{\lambda_m}||^2}\sqrt{1 - \left(\frac{\alpha_{m-1}}{\alpha_m - \alpha_{m-1}}\right)^2 \frac{||\hat{b}||^2}{C}} \\
&= \sqrt{1 - \left(\frac{\alpha_{m-1}}{\alpha_m - \alpha_{m-1}}\right)^2 \frac{||\hat{b}||^2}{C}}.
\end{aligned}$$

It is clear that as $C \to \infty$, the right-hand side of the inequality converges to 1. Since $1 \geq \rho\left(b^* - \hat{b}, \frac{\sqrt{C}}{||\hat{b}^{\lambda_m}||}\left[\sum_{l=m}^{n} \underline{\hat{b}}_l u^l\right]\right)$, applying the Squeeze Theorem yields the desired result. $\square$