Economic Theory
https://doi.org/10.1007/s00199-026-01739-8

RESEARCH ARTICLE

Endogenous network formation with diffusion incentives

Noémie Cabau[1]

Received: 29 March 2024 / Accepted: 6 July 2026

Abstract
This paper proposes a non-cooperative game of network formation in which creating access to the network – for oneself and others – is a public good. Link formation therefore reflects a trade-off between the social value of links and their private cost. The paper characterizes the architecture of strict Nash networks. Strict Nash equilibria are either flat or hierarchical. Flat equilibria consist of a wheel (i.e., a directed cycle) that may or may not encompass all agents. Hierarchical equilibria are trees in which some agents act as intermediaries between otherwise disconnected groups. When the cost of link formation is sufficiently low, the exhaustive wheel is the unique equilibrium architecture and best-response dynamics converge to it with probability one. For intermediate costs, multiple equilibrium architectures coexist and can be Pareto-ranked: the exhaustive wheels dominate the non-exhaustive ones and the empty network, and trees dominate the empty network.

Keywords Non-cooperative network formation · Intermediation · Non-empty disconnected networks

JEL Classifications C72 · D85

1 Introduction

In many models of network formation, links are valuable because they provide access to other individuals. Accordingly, agents decide which links to form by weighing the private benefits of reaching others against the cost of establishing and maintaining those links. This paper studies a different source of network value. Agents benefit not only from accessing others, but also from the contributions their links make to others' access. The paper asks how valuing this social dimension of one's links changes equilibrium network architectures.

This question relates to two strands of the literature. The first is the literature on network formation with private network benefits, initiated by Bala and Goyal (2000)

✉ Noémie Cabau
noemie.cabau@uni-corvinus.hu

[1] Institute of Finance, Corvinus University of Budapest, Fővám tér 8, 1093 Budapest, Hungary

and Jackson and Wolinsky (1996), where links are valuable because they grant access to other agents.[1] The second considers intermediation as a source of rival benefits that decrease with the number of intermediaries along a network connection (Goyal and Vega-Redondo 2007; Manea 2018; Siedlarek 2025; Choi et al. 2017; Kleinberg et al. 2008). By contrast, the present paper studies non-rival intermediation benefits: agents value their participation in network connections whenever they are essential to them.

This change in the source of network benefits, from private to social, has important implications. Equilibrium networks need not be connected. In particular, they may be nonempty and disconnected, with groups of agents that either do not interact at all or interact asymmetrically, in the sense that one group can access the other but not vice versa.

To study environments in which creating access is a public good, I develop the following non-cooperative game of network formation. A set of homogeneous agents simultaneously decides which costly links to form. Each link is paid for and maintained by the agent who initiates it. The resulting network determines not only whom each agent can access, but also the extent to which each agent facilitates the access of others.

More precisely, agent i benefits from mediating the connection between agents j and k whenever i is essential to that connection, that is, whenever every path from j to k passes through i. The specification of network benefits is deliberately general and may depend on network characteristics other than intermediation. For example, contributing to the connections of well-connected individuals may yield some prestige premium. Instead, agents may care about the relative number of connections they control. Individual decisions generally reflect a trade-off between the social value of links and the private cost of forming and maintaining them. The equilibrium concept is the strict Nash equilibrium, so that every player has a unique best response to the others' strategies.

The analysis considers both directed and undirected links. The distinction is important because it determines whether access created by a link is unilateral or reciprocal. With directed links, a link formed by agent i with agent j allows i to access every agent reachable from j, but it does not allow j to access the agents reachable from i. By contrast, with undirected links, a link formed by i and paid for by i allows both agents to access every individual reachable from either of them.

The paper completely characterizes the architecture of strict Nash networks. The characterization differs sharply depending on whether links are directed or undirected. With directed links, strict Nash networks are either flat or hierarchical. Flat architectures consist of a wheel (i.e., a directed cycle) that may or may not encompass all agents. In contrast, hierarchical architectures are trees of components in which only the root component may contain several agents. Figure 2 illustrates both cases. With undirected links, every nonempty strict Nash network is a center-sponsored star, illustrated on the right of Fig. 1.

[1] Other models have extended (Bala and Goyal 2000)'s setup. E.g., Galeotti et al. (2006) allows for heterogeneous players vis-à-vis the benefits and costs of forming links; Feri (2007) and Hojman and Szeidl (2008) analyze dynamic and static network formation, respectively; Olaizola and Valenciano (2018) proposes a unified framework encompassing both (Bala and Goyal 2000) and Jackson and Wolinsky (1996). Recent work has also studied endogenous network formation in collaborative R&D environments (Billand et al. 2026).

Fig. 1 A wheel on the left, a star on the right

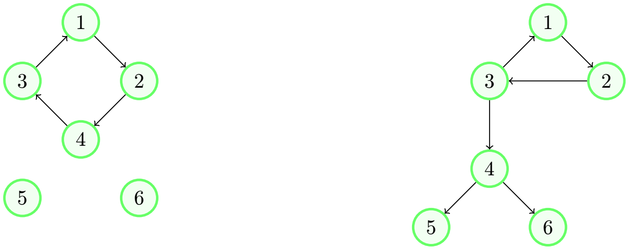

Fig. 2 Strict Nash networks: a non-exhaustive wheel on the left, a tree on the right

Three forces explain the architecture of strict Nash networks. The first two are economic and reflect the incentives created by the payoff function, whereas the third is purely a consequence of the strict Nash equilibrium refinement. First, in a Nash network, every maintained link must be essential to at least one network connection created by the agent who sponsors it. Otherwise, the agent could remove the link without affecting access, thereby saving its cost. Second, every maintained link must maximize the access it creates, given the rest of the network. Otherwise, the agent could redirect the link toward another part of the network, increase the access they provide to themselves and others, and thereby increase their payoff. Together, these two incentives discipline the architecture of Nash networks around a small number of essential intermediaries. Third, a strict Nash equilibrium requires every player to have a unique best response. The refinement rules out architectures in which an agent can replace one link with another without changing the set of network connections. As such, it selects the architectures with the fewest links. Consequently, groups of mutually connected agents form wheels under directed links and center-sponsored stars under undirected links.

These forces also govern the relations between components. If two components are not ordered hierarchically, one of them must be an isolated singleton. Otherwise, an agent in the smaller component could profitably redirect one of their links toward the larger component and thereby increase the access they provide. If two components are ordered hierarchically and links are directed, an agent whose component is accessed by the other can profitably redirect one of their links toward the component that accesses it. Doing so expands the access the agent provides and, consequently, their payoff. The uniqueness requirement imposed by the strict Nash equilibrium further rules out disconnected architectures in which an agent can replace one link with another without affecting the resulting network. Together, these restrictions imply that every strict Nash network has one component that either does not access anyone or accesses everyone outside it.

The characterization also yields simple existence conditions. In equilibrium, agents who maintain links must not benefit from severing existing links, while agents outside a component must not benefit from creating new links toward it. These conditions almost completely characterize the ranges of link-formation costs for which each equilibrium architecture exists.

The model applies to environments in which forming a well-connected group through a network benefits each individual – that is, the network generates a public good. Depending on the application, this public good may take the form of better information, greater participation, or improved access to valuable interactions. Individuals, therefore, choose links that enable as many network connections as possible.

One application concerns group communication and, more generally, situations in which individuals privately benefit from belonging to a well-informed group. Each agent contributes to the group's aggregate knowledge by collecting information through their network connections and sharing it with the rest of the group. For example, agents may be workers collaborating on a joint project, in which keeping teammates informed improves collective performance. Likewise, lenders may share information about borrowers' credit records (Padilla and Pagano 2000). By improving the information available to other lenders, such exchanges generate positive externalities by fostering better borrower discipline and reducing systemic risk.

Another application is diffusion auctions. Recent literature in computer science studies auctions in which sellers rely on participants to spread information about a sale to increase participation. Because informing additional buyers increases the auctioneer's expected revenue, it may be in the auctioneer's interest to reward participants for contributing to information diffusion. Li et al. (2022) propose an extension of the Vickrey-Clarke-Groves mechanism that is individually rational, incentive compatible, and efficient. Under that mechanism, each participant receives the increase in social welfare generated by their diffusion activity. This paper examines the networks that emerge when agents anticipate such diffusion rewards.

The remainder of the paper is organized as follows. Section 2 introduces the model and relates it to the diffusion auction application. Section 3 characterizes strict Nash networks and studies their dynamic stability. Section 4 extends the analysis to initial networks and undirected links, and compares the model with the network formation games of Bala and Goyal (2000) and Goyal and Vega-Redondo (2007). Section 5 concludes.

2 The environment

This section presents the model and discusses an application that motivates the payoff specification.

2.1 The model

A finite set of agents, $N = \{1, \ldots, n\}$, forms a network by creating links with one another. To rule out trivial cases, assume $n \geq 5$. Network formation is modeled as a non-cooperative game as follows. Agent i's strategy, g_i, is a set of agents toward whom i forms links. Links are *directed*; $j \in g_i$ is equivalent to $i \to j$, and means that "i can access j and all individuals reachable from j". The converse, however, is not true. Since an agent can form or not form links with each of the other $n - 1$ agents, the set of pure strategies of agent i is the power set of $N \setminus \{i\}$, $G_i = \mathcal{P}(N \setminus \{i\})$. Let $G = \Pi_{i \in N} G_i$ denote the space of pure strategy profiles. For simplicity, I consider only pure strategies. I use g to denote both the strategy profile $(g_1, \ldots, g_n)$ and the network formed by that profile. At times, a profile g is written as (g_i, g_{-i}), where g_{-i} denotes $(g_1, \ldots, g_{i-1}, g_{i+1}, \ldots, g_n)$.

Agent i's payoff depends on the cost of forming and maintaining links and on the benefits generated by their position in the network. If agent i chooses strategy g_i, they incur cost $C(|g_i|)$, where $|g_i|$ is the number of links formed by i. I assume that $C(0) = 0$, and that C is strictly increasing.

Agents benefit from their contribution to the reach of others. For any $j \in N$, agent j's *reach* is the number of individuals whom j can access in the network g. Agent j can access agent k if either $j \equiv k$ or there exists a path from j to k. A *path* from j to k is a sequence of links, $j \to i_1 \ldots \to i_k$, such that $i_k \equiv k$, $k \geq 1$, along which all agents are distinct. I denote by $r_j(g)$ agent j's reach in the network g; note that $r_j(g) \in \{1, \ldots, n\}$.

For any $(i, j) \in N \times N$, agent i's contribution to agent j's reach is defined as the decrease in j's reach when agent i leaves the network. Formally, the resulting network – denoted by $\tilde{g}^i$ – is defined on the set of nodes $N \setminus \{i\}$ and on the set of links $\{a \to b | a \in N \setminus \{i\}, b \in N \setminus \{i\}, b \in g_a\}$. Agent i's contribution to agent j's reach is denoted by $\Delta_{ij}(g)$ and is given by:

$$\Delta_{ij}(g) = r_j(g) - r_j(\tilde{g}^i).$$

For every $i \in N$, $r_i(g) = \Delta_{ii}(g)$. Define each agent's payoff function $u_i : G \to \mathbb{R}$ as additive-separable between benefits and costs:

$$u_i(g) = \sum_{j \in N} \Phi(r_j(g), \Delta_{ij}(g)) - C(|g_i|). \tag{1}$$

In the above, Φ is the benefit that agent i earns from their contribution to the network. It satisfies the following properties.

Assumption 1 Consider the class of functions $\Phi : \mathbb{Z}_+^2 \to \mathbb{R}_+$, where $\mathbb{Z}_+$ denotes the set of non-negative integers, and $\mathbb{R}_+$ the set of non-negative reals, such that:

1. $\Phi(., 0) = 0$ and $\Phi(., y) > 0$ for all $y > 0$.
2. $\Phi(x, y) < \Phi(x, y + 1)$, for all $x > y$, for all y.
3. $\Phi(x + 1, y + 1) > \Phi(x, y)$ for all $x \geq y$.

The first statement implies that an agent earns a benefit only if they make a positive contribution to some individual's reach. The second statement captures the value of intermediation. The third statement implies that agents benefit from increases in others' reach when those increases are entirely attributable to them.

I assume that agents choose their strategies simultaneously. A network g is *Nash* if, for all $i \in N$ and all $h_i \in G_i$, $u_i(g_i, g_{-i}) \geq u_i(h_i, g_{-i})$; each agent's strategy is a best response to the strategies of the others. A network g is *strict Nash* if, for all $i \in N$ and all $h_i \in G_i \setminus \{g_i\}$, $u_i(g_i, g_{-i}) > u_i(h_i, g_{-i})$; g_i is the unique best response of i to g_{-i}.

Before undertaking the analysis of the model, it is useful to introduce some terminology. A set $A \subseteq N$ is called a *component* if, for every ordered pair of agents $(i, j) \in A \times A$, i can reach j and there is no strict superset A' of A for which this is true. A component A is a *singleton* if it has a unique element, i.e., $|A| = 1$. A component A is *minimally connected* if A is no longer a component upon deletion of a link $i \to j$ between two agents i and j in A. A component A forms a *wheel* when the agents in A can be arranged as $\{i_1, \ldots, i_k\}$ for some $k \in \{2, \ldots, n\}$, $i_1 \to i_2 \ldots \to i_k \to i_1$, and there are no other links between the agents in A. A network is *connected* if it has one component. A network that is not connected is *disconnected*. The *empty network* is the network with no links.

2.2 Diffusion auctions

The payoff specification introduced above admits several interpretations. This subsection presents a diffusion auction as a motivating example. The objective is not to derive the equilibrium results that follow, but rather to illustrate an environment in which agents benefit from making themselves and others accessible. The equilibrium analysis in the remainder of the paper applies to any payoff satisfying Assumption 1.

Suppose that network formation precedes an auction in which the set N represents potential buyers. Agents have independent private values for the item on sale, drawn from the uniform distribution on the unit interval. The auctioneer cannot publicly advertise the sale. Instead, she initially informs one randomly selected agent in the network, called the *information seed* and denoted j_0. Only the agents reachable from j_0 can learn about the sale and participate.

The key observation is that some agents are essential for information diffusion. By maintaining links, these agents allow additional buyers to learn about the auction. A sale mechanism may therefore reward agents according to the contribution they make to others' access. Consider the following extension of the Vickrey–Clarke–Groves mechanism.

Every informed agent i decides to participate or not; if they do, they report some $t'_i = (v'_i, g'_i)$ to the auctioneer, where $v'_i \in [0, 1]$ is the reported valuation and $g'_i \in G_i$ the reported links. The item is allocated to the participant with the highest reported valuation. Let $P(j_0; g')$ denote the set of participants when j_0 is the information seed and g' is the reported network, and define the highest reported valuation as

$$v^*(g') = \max_{j \in P(j_0; g')} v'_j.$$

Fig. 3 A line network

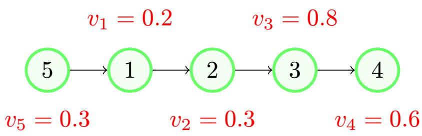

Agent i receives the transfer:

$$\pi_i(t') = \begin{cases} 0 & \text{if } i \notin P(j_0; g'), \\ v^*(g') - v^*(\tilde{g}'^i) & \text{if } i \in P(j_0; g'), \end{cases}$$

where $\tilde{g}'^i$ denotes the reported network after removing agent i and the links formed by i. Thus, an agent is rewarded according to the reduction in welfare caused by their removal. An agent receives a positive transfer only if their removal lowers the highest valuation among the remaining participants. The transfer to agent i increases with the highest valuation among the participants who are reachable only through i.

The mechanism is individually rational and incentive compatible (see Li et al. (2022)). Figure 3 illustrates its working. Suppose that agent 1 is the information seed. Agents 1 to 4 participate, whereas agent 5 remains uninformed because it is unreachable from agent 1. Agent 3 wins the auction since $v_3 = 0.8$ is the highest reported valuation among participants. Agents 1, 2, and 3 are all essential for connecting the information seed to the winner. Removing agent 3 leaves agent 2 as the highest-valued reachable participant, so agent 3 pays $v_2 = 0.3$. Removing agent 2 would result in allocating the item to agent 1, generating a welfare loss of $v_3 - v_1 = 0.6$. Finally, removing agent 1 prevents the sale altogether, generating a welfare loss of $v_3 = 0.8$. Hence, agents 1 and 2 receive positive transfers because their participation enables access to buyers with higher valuations.

Fix a network g. The interim expected payoff of agent i, after network formation but before the information seed is chosen and transfers are realized, is:

$$u_i(g) = \frac{1}{n} \sum_{j_0 \in N} \left(\frac{r_{j_0}(g)}{r_{j_0}(g) + 1} - \left(\frac{r_{j_0}(g) - \Delta_{ij_0}(g)}{r_{j_0}(g) - \Delta_{ij_0}(g) + 1} \right) \right) - C(|g_i|).$$

The interpretation is as follows. When j_0 is the information seed, $r_{j_0}(g)$ buyers participate in the auction. Removing agent i reduces that number to $r_{j_0}(g) - \Delta_{ij_0}(g)$. Since valuations are uniformly distributed on $[0, 1]$, the expected highest valuation among m buyers is $m/(m+1)$. The term in brackets, therefore, measures the expected welfare loss generated by removing agent i. Averaging across all possible information seeds yields agent i's expected network benefit.

The diffusion auction thus provides one environment in which agents benefit from their contribution to others' access. It is established in Appendix A that the resulting payoff satisfies Assumption 1. The Appendix also shows that introducing a reserve price guarantees budget balance when the number of agents is finite.

3 Equilibrium networks

The section features a tight characterization of equilibrium networks and provides conditions for their existence. Unless stated otherwise, all proofs are relegated to the Appendix.

3.1 Properties of Nash networks

The payoff specification rewards agents for their contribution to making the network accessible to themselves and others. Consequently, a best response needs to satisfy two properties. First, every maintained link must be essential to some network connection created by the agent who sponsors it. Otherwise, the agent could remove the link without affecting access, thereby saving its cost. Second, every maintained link must maximize the access it creates, given the other links in the network. Otherwise, the agent could redirect it toward another part of the network and obtain a strictly higher payoff. Proposition 1 shows that these two incentives impose a strong restriction on links pointing toward the same agent.

Proposition 1 *If g is Nash and there are three distinct players i, j, k such that $k \in g_i$ and $k \in g_j$, then i, j and k belong to the same component. Also, all paths from i to j pass through k, and symmetrically, all paths from j to i pass through k.*

Proof In some network g, assume that there are three agents i, j and k such that $k \in g_i$, $k \in g_j$, and they do not belong to the same component. Assume first that i and j are in two different components. Without loss of generality, suppose i does not have a path to j. Let i deviate to $g_i' = (g_i \setminus \{k\}) \cup \{j\}$: i cuts the link with k and simultaneously adds a link towards j. Let $g' = (g_i', g_{-i})$ denote the resulting network. The deviation preserves access to k via the link $i \rightarrow j$ while creating access to the component of agent j. Thus, for any agent $a \in N$ such that $r_a(g') \neq r_a(g)$, it holds that $r_a(g') - r_a(g) = \Delta_{ia}(g') - \Delta_{ia}(g) > 0$. Such an agent a exists (as i is one of them). Since $|g_i'| = |g_i|$, we conclude that g_i' is a strictly profitable deviation for i. Thus, g is not Nash.

Suppose instead that i, j are in the same component, A. Without loss of generality, suppose there is a path from i to j that does not pass through k. Let i deviate to $g_i' = g_i \setminus \{k\}$, and let $g' = (g_i', g_{-i})$ be the resulting network. Agent i still has a path to j in g' (those that do not pass through k). Since j maintains the link $j \rightarrow k$, there is still a path from i to k in g'. Hence, the link $i \rightarrow k$ is made redundant by the link $j \rightarrow k$. Therefore, $r_a(g) = r_a(g')$ and $\Delta_{ia}(g) = \Delta_{ia}(g')$ for all $a \in N$. Since $|g_i'| < |g_i|$, g_i' is a strictly profitable deviation for i. Thus, g is not Nash. □

Proposition 1 is the cornerstone of the equilibrium characterization. It implies that whenever two agents form links toward the same individual, that individual must be essential for their mutual connection. In equilibrium, links therefore cannot duplicate existing paths through the network. Instead, they organize access around a small number of essential "bottlenecks" such as agent k in the Proposition. The remainder of this section develops the structural implications of this observation.

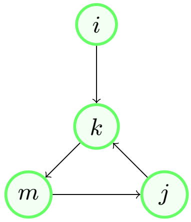

Fig. 4 A network that is not Nash, where $A = \{i\}$ and $B = \{j, k, m\}$. Agent j could strictly improve their payoff by cutting the link with k and adding a link with i

Corollary 2 develops the implications on the pattern of access between components. To illustrate this idea, I consider the relation $\hookrightarrow$ over the set of components of a network: $A \hookrightarrow B$ when component A accesses B (i.e., all agents in A can reach all agents in B). A is said to be *maximal* (*minimal*) if there is no component $B \neq A$ such that $B \hookrightarrow A$ ($A \hookrightarrow B$). A is said to be *isolated* if it is both maximal and minimal.

Corollary 2 *In a Nash network, (1) A component that is not maximal is a singleton, (2) A component is either a singleton or it is minimally connected, (3) If components A and B are maximal and $A \neq B$, A or B is an isolated singleton.*

Each statement can be intuitively understood. For Statement 1, consider a network with two components A and B such that $A \hookrightarrow B$ but not the other way around. Graphically, there is a bridging link from A to B such as $i \rightarrow k$ in Fig. 4. If, unlike stated in the Corollary, B has several agents, then one of them has a link with k (see j in the Figure). This contradicts Proposition 1 as i and j are in different components. For Statement 2, observe that a non-minimally connected component has two agents connected by a redundant link. As previously mentioned, the agent who maintains it could spare its cost and improve their payoff.

Statement 3 follows from the earlier observation that each link in a Nash network must maximize the access created by that link, given the other links in the network. Suppose, contrary to the statement, that there are two distinct maximal components, A and B, neither of which is an isolated singleton. Since A is maximal, a deviation by one of its members affects only the reach of the agents in A. Suppose that the agents in B have a weakly higher reach than those in A. Then an agent in A can cut all their existing links and instead form a single link to an agent in B. This deviation increases the reach of every agent in A without affecting the reach of any agent outside A. Since the deviating agent forms only one link, the deviation does not increase link costs, while it strictly increases the agent's network benefit. Hence, the deviation is strictly profitable.

Corollary 2 provides sharp predictions about the overall structure of Nash networks. Either they are *flat*, with no comparable components. The empty network is one of them; the non-empty ones have a unique minimally connected component surrounded by isolated singletons (Statement 3 in Corollary 2). Alternatively, a Nash network can be *hierarchical*, with one component accessing all others. It may help to think of hierarchical networks as trees of components. The root component is the greatest element of $\hookrightarrow$, and it is either a singleton or minimally connected (by Statement 2 in Corollary 2). The rest of the components are singletons (by Statement 1 in Corollary 2), and those without links are called *leaves*. Finally, a Nash network may combine flat and hierarchical features, with one tree of components surrounded by isolated singletons (by Statement 3 in Corollary 2).

Statement 2 in Corollary 2 states that a component other than a singleton must be minimally connected. Many architectures satisfy this property: the star network, Dutch Windmills, and overlapping cycles are examples. Consequently, the Nash concept alone leaves substantial indeterminacy regarding the architecture of equilibrium components. The next section shows that, holding the set of network connections fixed, the strict Nash refinement selects the architecture with the fewest links.

3.2 Characterization and existence of strict Nash networks

The strict Nash equilibrium further requires that each player has a unique best response to the others' equilibrium strategies. As shown below, every nonempty strict Nash network has one wheel that either does not access anyone or accesses everyone outside it. The argument proceeds in three steps.

First, in strict Nash networks, a component with several agents forms a wheel. Indeed, by Proposition 1, two agents i and j who have a link with the same individual k each have a payoff-equivalent strategy (e.g., i maintains all of their links, except for cutting the link with k and adding one to j); see Figure 5 for an illustration.

Second, strict Nash networks are either flat or hierarchical and do not exhibit both features. This is because a player in a tree who has initiated a link with a leaf would be equally well off by cutting it and adding one to an isolated singleton. Such indifference is incompatible with strictness.

Last, the root component of a hierarchical strict Nash network must encompass several agents and, therefore, forms a wheel. To see why, consider a tree whose root component is a singleton. Notice that the links initiated by the root do not confer it any intermediation benefits, since the root is the only agent that uses them. Consider a leaf of the tree; by forming a single link to the root, the former would also access the rest of the tree. Furthermore, the leaf gains additional intermediation benefits from these agents (if any) between them and the root. Together, these observations imply that whenever the root prefers maintaining its links to cutting them, a leaf prefers linking to the root. Therefore, the two agents cannot both play a strict best response, and the network is not a strict Nash equilibrium.

Proposition 3 offers a categorization of strict Nash networks based on their architecture. Its content can be mostly recovered from the previous three observations: Flat and nonempty strict Nash networks have one wheel that may or may not encompass all individuals; Hierarchical strict Nash networks have a wheel as root component, and the rest of the agents form "branches" that grant the wheel access to every agent outside it. See Table 1 for definitions and further details.

Fig. 5 For agent 4, $g_4 = \{1\}$ and $g'_4 = \{3\}$ are payoff-equivalent strategies

Proposition 3 *If g is a strict Nash network, then g has one of the following architectures: (1) the empty network, (2) a non-exhaustive wheel, (3) an exhaustive wheel, (4) a root-leaves tree, (5) a tree with intermediaries.*

Proposition 4 provides necessary and sufficient conditions for the existence of the first four architectures listed in Proposition 3. I exclude trees with intermediaries because their heterogeneity makes it difficult to derive network-specific conditions.

The result is organized around three types of deviations from an agent's current strategy: exclusively adding links; exclusively cutting links; last, first rewiring the existing links, then adding or removing links from that set. The Proposition compares the payoff of each agent to that from their *most profitable* deviation, this, for each architecture. For each agent in all architectures except the non-exhaustive wheels, the most profitable deviations involve exclusively adding or cutting links. The case of the non-exhaustive wheels is longer because, for a wheel member, the third type of deviation may be the most profitable (cutting the link in the wheel, then adding links to isolated singletons).

Proposition 4 *(1) The empty network is strict Nash if and only if no agent wants to form links: $\Phi(1, 1) > \Phi(x + 1, x + 1) - C(x)$ for all $x \in \{1, \dots, n - 1\}$. (2) An exhaustive wheel is strict Nash if and only if no agent wants to cut their link: $\sum_{k=1}^{n} \Phi(n, k) - C(1) > \sum_{k=1}^{n} \Phi(k, 1)$. (3) A non-exhaustive wheel with w agents is strict Nash if and only if:*

- *(3.a) No isolated singleton gains from forming one link to the wheel and either none or some links to other isolated singletons: $\Phi(1, 1) > \Phi(w + x, w + x) - C(x)$, $\forall x \in \{1, \dots, n - w\}$.*
- *(3.b) No wheel member gains from:*
  - *(3.b.i) cutting their link in the wheel: $\sum_{k=1}^{w} \Phi(w, k) - C(1) > \sum_{k=1}^{w} \Phi(k, 1)$,*
  - *(3.b.ii) cutting their link in the wheel and forming links to isolated singletons: $\sum_{k=1}^{w} \Phi(w, k) - C(1) > \sum_{k=x+1}^{w+x} \Phi(k, x + 1) - C(x)$, $\forall x \in \{1, \dots, n - w\}$,*
  - *(3.b.iii) maintaining their link in the wheel and adding some to isolated singletons: $\sum_{k=1}^{w} \Phi(w, k) - C(1) > \sum_{k=x}^{w+x-1} \Phi(w + x - 1, k) - C(x)$, $\forall x \in \{2, \dots, n - w + 1\}$.*

(4) A root-leaves tree g whose root component is denoted by A is strict Nash if and only if:[2]

- *(4.a) A leaf does not gain from forming one link to a root member: $\Phi(1, 1) > \Phi(n, n) - C(1)$.*
- *(4.b) Any root member i does not gain from:*
  - *(4.b.i) cutting all their links: $\sum_{j \in A} \Phi(n, \Delta_{ij}(g)) - C(|g_i|) > \sum_{j \in A} \Phi(n - \Delta_{ij}(g) + 1, 1)$,*
  - *(4.b.ii) cutting links with some leaves (if $|g_i| > 1$): $\sum_{j \in A} \Phi(n, \Delta_{ij}(g)) - C(|g_i|) > \sum_{j \in A} \Phi(n - x, \Delta_{ij}(g) - x) - C(|g_i| - x)$, $\forall x \in \{1, \dots, |g_i| - 1\}$.*

[2] Note, if the root component consists of the agents indexed from 1 to w and the links between the root members are $1 \to 2 \dots \to w \to 1$, $\Delta_{ij}(g) = 1_{j>i} \sum_{p=i}^{j-1} |g_p| + \gtrdot 1_{j<i} \left(\sum_{p=i}^{w} |g_p| + \sum_{p=1}^{j-1} |g_p|\right)$, for any $i, j \in A$.

Table 1 Strict Nash networks, $n = 5$

Flat Architecture	Properties	Hierarchical Architecture	Properties
Empty network 	Each agent is isolated, with no link adjacent to them.	*Root-leaves trees* 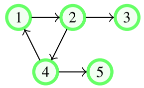	A directed rooted tree of components encompassing all individuals. One wheel component, the root, encompasses 2 to $n-1$ individuals. Each other component is a singleton. All singletons are leaves.
Non-exhaustive wheels 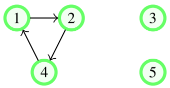	One wheel component encompassing from 3 to $n-1$ individuals. Individuals outside the wheel are isolated, with no adjacent links.	*Trees with intermediaries* 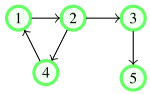	A directed rooted tree of components encompassing all individuals. One wheel component, the root, encompasses 2 to $n-2$ individuals. Each other component is a singleton. Some singletons maintain links and serve as intermediaries.
Exhaustive wheels 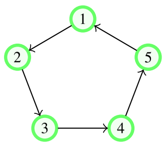	One wheel component encompassing all n individuals.		

The conditions in Proposition 4 highlight three incentives. First, some inequalities rule out *profitable link formation*: agents should not gain from expanding access by adding new links. Second, some inequalities rule out *profitable link severance*: the contribution generated by existing links must exceed the savings from deleting them. Finally, Condition (3.b.ii) rules out *profitable "rewiring" deviations* in non-exhaustive wheels, whereby a wheel member cuts the link in the wheel and instead forms links with isolated singletons. The Proposition focuses on the most profitable deviations within each architecture. For example, Statements (3.a) and (4.a) are restricted to strategies for a singleton that include a link to the wheel component of their network. This is because such a link generates more access than a link to another singleton, ceteris paribus. Similarly, in Statement (4.b.ii), a root member who wishes to cut a link prefers to cut one with a leaf rather than one in the wheel. For non-exhaustive wheels, additional deviations must be considered. In particular, Statement (3.b.ii) compares a wheel member's payoff with that obtained from cutting the link in the wheel and forming links with isolated singletons, thereby transforming the wheel into a directed tree.

The conditions in Proposition 4 imply a Pareto ranking of equilibrium architectures. Architectures that generate more opportunities for access and intermediation remain sustainable for higher linking costs and yield weakly higher payoffs whenever they exist.

Corollary 5 *(1) If a non-exhaustive wheel or a tree (root-leaves or with intermediaries) is strict Nash, so are the exhaustive wheels and the empty network. (2) When a non-exhaustive wheel is strict Nash, the exhaustive wheels payoff-dominate it. Also, when a tree is strict Nash, the tree payoff dominates the empty network.*

Corollary 5 can be summarized as follows. When the cost of link formation is low so that forming no links is a strictly dominated strategy, the exhaustive wheel is the unique strict Nash architecture. Conversely, when the cost is high so that a link in an exhaustive wheel is not worth maintaining, the empty network is the unique equilibrium. For intermediate values of the cost of link formation, non-exhaustive wheels and trees may also arise in equilibrium. However, in such cases, exhaustive wheels and trees payoff-dominate some of the other equilibria.

The intuition is that exhaustive wheels and trees generate more opportunities for intermediation. In an exhaustive wheel, wheel members bear the same link costs as in a non-exhaustive wheel but contribute to the reach of more agents. Likewise, even a leaf without links in a tree benefits from being accessible to other agents, whereas an agent in the empty network does not.

3.3 Dynamic stability

Given the extent of the coordination problem in this game, one may wonder whether agents can reach an equilibrium through decentralized adjustments. The next result addresses this question for sufficiently low link-formation costs, where the exhaustive wheel is the unique strict Nash architecture.

Fig. 6 Cycling of best response dynamics

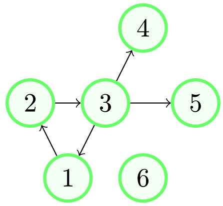

I study a simple best-response dynamic in which agents revise their linking decisions sequentially over time. Consider a dynamic process in discrete time where the state at time t is the prevailing network $g(t) = (g_1(t), \dots, g_n(t))$, for each $t = 0, 1, \dots$. In each period $t > 0$, one player is randomly chosen and given the opportunity to play a best response to the profile $g_{-i}(t)$ played by the others. If player i has several best responses, each of them is assumed to be chosen with positive probability.

The next Proposition establishes that when the cost of link formation is low so that forming no links is a weakly dominated strategy, the dynamic process converges to an exhaustive wheel.

Proposition 6 *If* $\Phi(x+1, x+1) - C(x) \geq \Phi(1, 1)$ *for some* $x \in \{1, \dots, n-1\}$, *every strict Nash network is an exhaustive wheel, and the dynamic process converges to an exhaustive wheel with probability* 1.

Intuitively, the heart of the matter is to prove that the dynamics do not cycle. Concretely, cycling may occur when the dynamic process reaches a Nash equilibrium that is not strict and for which, whenever an agent has several best responses, all of them preserve the architecture of the current network. An example of this is presented in Fig. 6. The network satisfies Proposition 1, hence, it is a Nash candidate. However, it is never strict Nash as agent 3 has several payoff-equivalent strategies: the current one $g_3 = \{1, 4, 5\}$, but also $g_3' = \{1, 5, 6\}$ and $g_3'' = \{1, 4, 6\}$. Suppose the network is Nash, agent 3's best replies are those listed above, and each other player is playing their unique best response. Then, all networks that subsequently occur with positive probability have that architecture; only the identity of the leaves changes. E.g., there is a positive probability that the leaves in $g(t+1)$ are agents 5 and 6, and again, agents 4 and 5 in $g(t+2)$.

The condition on the cost of link formation implies that an agent without links weakly gains from forming some. Therefore, any architecture with isolated agents, such as in Fig. 6, cannot persist under the best-response dynamics, and the process moves away from it with positive probability.

4 Discussions

This section provides two extensions of the game. The first accommodates the existence of some original network; the second considers the case where links are undirected. The last subsection compares the current model with two closely related games by Bala and Goyal (2000) and Goyal and Vega-Redondo (2007).

4.1 Existence of an initial network

In the baseline model, the link formation game starts from the empty network. In many applications, however, agents may inherit pre-existing relationships before the link formation stage (Joshi et al. 2020). To capture this possibility, suppose that some network g_0 exists prior to the agents' decisions. For simplicity, assume that g_0 consists of one component – denoted by A_0 – with several agents, and a collection of isolated singletons. The links in g_0 cannot be severed. The strategy space and individual payoffs are otherwise unchanged, and the equilibrium concept remains the strict Nash equilibrium.

The presence of the component A_0 modifies the set of admissible equilibrium architectures, but the logic developed in Section 3.2 largely carries over. A key observation is that in a strict Nash network, no agent outside A_0 forms a link toward A_0. Indeed, linking to any member of A_0 grants access to the entire component, so links to different members of A_0 yield identical benefits. As a result, any agent outside A_0 who forms a link to a member of A_0 has a payoff-equivalent strategy obtained by redirecting that link to another member of A_0, which is incompatible with strict Nash equilibrium. Moreover, links between members of A_0 generate no additional benefit, since these agents are already connected through the initial network.

The equilibrium networks are close to those identified in the baseline model. Agents outside A_0 may form a wheel component among themselves. However, a wheel member can always replace a link within the wheel with a link towards A_0. Therefore, this architecture is sustainable in a strict Nash equilibrium only if the wheel contains more agents than A_0, so that replacing a wheel link by a link towards A_0 would reduce the agent's access and, consequently, the benefits generated by that link. Likewise, agents may form a hierarchical architecture rooted at A_0. Since no agent can link toward A_0 in a strict Nash equilibrium, the root component is fixed, and the remaining components are singletons (by Corollary 2 Statement 1).

The next Proposition summarizes the possible strict Nash architectures in the presence of an initial network and uses the following terminology. An *exhaustive tree* is a directed rooted tree of components that encompasses all agents; the root is A_0, and all other components are singletons. An *augmented non-exhaustive wheel* is a flat network in which no two components are comparable via $\hookrightarrow$; one component is A_0, another component A_1 forms a wheel with more agents than A_0 (i.e., $|A_1| > |A_0|$), and the remaining components are isolated singletons.

Proposition 7 *Suppose that the link formation game begins from some initial network g_0, where g_0 is a flat network with one component A_0 containing several agents (i.e., $|A_0| \geq 2$) and the rest are isolated singletons. A network g is strict Nash only if it has one of the following architectures: (1) g_0, (2) an exhaustive tree, (3) an augmented non-exhaustive wheel.*

4.2 Undirected links

The baseline model assumes that links are directed: an agent who forms a link with some individual gains access to the agents reachable from the latter, but the converse

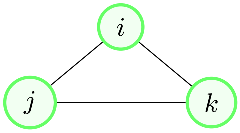

Fig. 7 Undirected cycle

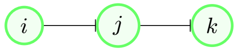

Fig. 8 Two consecutive links formed by different players

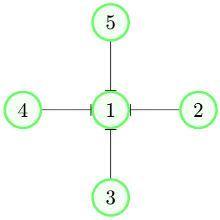

Fig. 9 A center-sponsored star

need not hold. In many environments, however, access is reciprocal. This subsection, therefore, considers the version of the game in which links are undirected. As will become clear, equilibrium predictions differ markedly. Whereas directed links support wheel and tree architectures, strict Nash equilibrium under undirected links is much more restrictive and selects either a center-sponsored star or the empty network.

With undirected links, agent i can reach agent j if either $i \equiv j$ or $i \neq j$ and there exists an *undirected path* $i - i_1 - \ldots - i_k$ such that $i_k = j$, $k \geq 1$, along which all agents are distinct. Therefore, in this version of the game, the relation $j \in g_i$ reads as "i forms a link with j, i pays for it; i can access all agents reachable from j and j can access all agents reachable from i". The strategy space is unchanged, and individual payoffs are still given by (1).

The characterization of strict Nash networks proceeds in three steps.

1. First, a (strict) Nash network is *acyclic*. To illustrate the argument concretely, consider the network in Fig. 7. Without loss of generality, suppose that agent i pays for the link $i - j$. Whether this link is maintained or cut does not affect the contribution that i makes to the others' reach. Hence, i obtains the same benefit in both cases. Since links are costly, i strictly prefers removing the link to save its cost.

2. Second, strict Nash equilibrium imposes a much stronger requirement than acyclicity. While every component must be a tree, no two consecutive links in that tree can be sponsored by different agents.
   To see why, suppose instead that there is a sequence of two consecutive links initiated by different players, as shown in Figs. 8 and 9. For clarity, a dash lies near the agent who paid for the link. Player k paid for the link with j, whereas j paid for the link with i. Because links are undirected, k's contribution to the others' reach is the same regardless of whether k links to i or to j. Agent k therefore has two payoff-equivalent strategies, so the network cannot be strict Nash.

It follows that the same agent must sponsor all links in a component. Therefore, any component with several agents is a *center-sponsored star*.

3. Finally, a strict Nash network is either connected or empty.
   The argument is twofold. First, there can be at most one center-sponsored star. If two such stars existed, the center of either star would strictly benefit from replacing one of its links to a peripheral agent by a link to an agent in the other star. The former grants access to a single individual, whereas the latter grants access to an entire component.
   Second, a strict Nash network containing a center-sponsored star cannot contain isolated singletons. From the center's perspective, replacing a peripheral agent in the star by an isolated singleton leaves the architecture unchanged: in both cases, the center allows access to exactly one additional agent. The center, therefore, has two payoff-equivalent strategies, so the network cannot be a strict Nash equilibrium.

In conclusion, a strict Nash network is either a center-sponsored star or the empty network. The remaining question is for which ranges of cost can each architecture be a strict Nash equilibrium. To answer this question, define

$$\gamma(x) = \Phi(x+1, x+1) + x\Phi(x+1, x) - C(x),$$

for $x \in \{0, \ldots, n-1\}$. It may help the reader to think of $\gamma(x)$ as the payoff to the central agent of a center-sponsored star with $(x+1)$ agents.

The center-sponsored star and the empty network cannot co-exist in equilibrium. Indeed, one obtains from the other through a unilateral deviation. From the empty network, an agent can form links with all other agents and thereby create a center-sponsored star; from a center-sponsored star, the center creates an empty network by cutting all its links.

Proposition 8 *Consider the function $\gamma(x)$ defined above. (1) If $\gamma(0) > \gamma(x)$ for all $x \in \{1, \ldots, n-1\}$, the empty network is the unique strict Nash network. (2) If $\gamma(n-1) > \gamma(x)$ for all $x \in \{0, \ldots, n-2\}$, a strict Nash network is a center-sponsored star. (3) Otherwise, $\gamma(0) \leq \gamma(x)$ for some $x > 0$ and $\gamma(n-1) \leq \gamma(x')$ for some $x' < n-1$, and there are no strict Nash networks.*

Unlike in the directed model, strict Nash equilibria need not exist under undirected links. Since the only candidate strict Nash architectures are the empty network and the center-sponsored star, a strict Nash equilibrium exists only if either 0 or $n-1$ is the unique global maximizer of γ. Whenever this condition fails, no strict Nash equilibrium exists.

4.3 Related network formation games

This section compares the equilibrium predictions of the current game with those of Bala and Goyal (2000) and Goyal and Vega-Redondo (2007). The key difference between the current model and the game of Bala and Goyal (2000) is that agents in their model evaluate links solely through their effect on their own reach. In contrast, agents in the present model also value the contribution their links make to others' reach. As shown below, this distinction has important implications for the set of equilibrium architectures.

In Bala and Goyal (2000), individual decisions in link formation depend on the trade-off between the *private* benefit of links, measured by reach, and the cost of link formation. For any agent $i \in N$, Bala and Goyal (2000)'s payoff function is

$$u_i(g) = \Psi(r_i(g), |g_i|),$$

where $r_i(g)$ is agent i's reach in network g, and $|g_i|$ is the number of links formed by i. The function Ψ is increasing in reach and decreasing in the number of links formed. To ease the comparison between the two models, consider the following specification of the above payoff:

$$u_i(g) = \Phi(r_i(g), r_i(g)) - C(|g_i|). \tag{2}$$

Recall that, for any network g and any agent i, $\Delta_{ii}(g) = r_i(g)$.

Fixing the network formed by the other players, the best-response set of a player may differ across the two games. The reason is that agents in Bala and Goyal (2000)'s model do not internalize the positive externalities that their links generate for others. To illustrate the difference, consider the non-exhaustive wheel in Fig. 10.

In Bala and Goyal (2000)'s game, agent 5 can tap higher benefits than any wheel member by mimicking their strategy. Indeed, if agent 5 adds a link to agent 1, the payoff from the deviation is $\Phi(5, 5) - C(1)$, which exceeds agent 3's current payoff of $\Phi(4, 4) - C(1)$. If agent 3 cuts the link with agent 1, the payoff from the deviation is $\Phi(1, 1)$, which is precisely agent 5's current payoff. Therefore, agents 3 and 5 cannot simultaneously be playing best responses.

The crucial difference in the current game is that agent 3's link generates value not only through the access it provides to agent 3, but also through the access it provides to the other members of the wheel. These intermediation benefits are absent from Bala and Goyal (2000)'s game. As a result, agents 3 and 5 may simultaneously be playing best responses in the current model. The same logic extends to trees. In Bala and Goyal (2000)'s game, a leaf can mimic the strategy of a root member and obtain the same reach, therefore, the same payoff. In contrast, root members in the current game receive additional intermediation benefits that may not be available to leaves.

Fig. 10 A non-exhaustive wheel

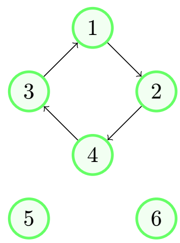

More generally, link formation in the current game reflects a trade-off between the social value of links and their private cost, whereas in Bala and Goyal (2000)'s model the trade-off is between the private value of links and their cost. Since the social value of a link weakly exceeds its private value, non-empty networks are equilibria for a wider range of costs in the current game. This also explains why the condition for the empty network coincides across the two games under specification (2): starting from the empty network, the social value and the private value of a newly formed link are identical.

The discussion above implies that the set of equilibrium architectures is larger in the current game. The following Corollary summarizes the comparison for directed links. Bala and Goyal (2000) show that a strict Nash network is either an exhaustive wheel or the empty network. By contrast, the current game admits additional equilibrium architectures, namely non-exhaustive wheels and trees.

Corollary 9 *(See Proposition 3.2 pp. 1195 in Bala and Goyal (2000).) Consider the payoffs in* (2) *and* (1)*. (1) If* $\Phi(n, n) - C(1) > \Phi(1, 1)$*, the exhaustive wheels are strict Nash in both games. (2) Otherwise, the empty network is the unique strict Nash network in Bala and Goyal (2000)'s game. In contrast, some architectures among the exhaustive wheels, the non-exhaustive wheels and the trees may be strict Nash in this game.*

Equilibrium predictions converge under undirected links. The arguments developed in Section 4.2 continue to apply because they do not depend on whether agents value the contribution that their links make to the reach of others. The underlying reason is the inherent symmetry in access afforded by undirected links: once an agent connects to a component, the particular individual through whom access is obtained becomes irrelevant. Consequently, both games admit the same strict Nash architectures under undirected links, namely the empty network and the center-sponsored star. The question of existence, therefore, drives the comparison. In the empty network, an agent who deviates and forms links becomes, in effect, the center of a center-sponsored star. Since links in a star generate positive externalities, the social value of a link exceeds its private value. Consequently, agents have a stronger incentive to form stars in the current game than in Bala and Goyal (2000)'s model. The empty network is therefore a strict Nash equilibrium for a wider range of costs in Bala and Goyal's game.

Corollary 10 *(See Proposition 4.2 pp. 1204 in Bala and Goyal (2000).) Consider the payoffs in* (2) *and* (1)*. (1) If* $\Phi(n, n) - C(n-1) > \Phi(x+1, x+1) - C(x)$ *for all* $x \in \{0, \ldots, n-2\}$*, a center-sponsored star is the unique strict Nash network in both games. (2) Otherwise, and if for some* $x \in \{1, \ldots, n-1\}$*,*

$$\Phi(x+1, x+1) + x\Phi(x+1, x) - C(x) > \Phi(1, 1),$$

a strict Nash network is a center-sponsored star in this game, and it is the empty network in Bala and Goyal (2000)'s game. (3) If $\Phi(1, 1) > \Phi(x+1, x+1) + x\Phi(x+1, x) - C(x)$ *for all* $x \in \{1, \ldots, n-1\}$*, the empty network is the unique strict Nash network in both games.*

Another closely related game is the bilateral trade model of Goyal and Vega-Redondo (2007). The comparison with their model can be viewed as a comparative-statics exercise on the nature of intermediation benefits. To isolate this dimension, I adopt the link-formation protocol and equilibrium concept of Goyal and Vega-Redondo (2007). Namely, links form cooperatively, benefits flow symmetrically between connected agents (i.e., links are undirected), and the equilibrium concept is strict bilateral equilibrium.

This way, the difference between the two models reduces to the nature of the benefits. In Goyal and Vega-Redondo (2007), intermediation benefits are *rival*: the value generated by a network connection is shared among all agents who are essential for it. In the current model, benefits are *non-rival*: an agent's reward from being essential in a connection does not depend on the number of other essential agents. The remainder discusses how the incentive to form links then changes, and whether the difference alters the set of equilibrium architectures.

As shown below, the answer is largely negative. When the number of agents is sufficiently large, both models select the same equilibrium architectures, namely the star network and the empty network. The forces sustaining these architectures, however, differ markedly.

I will first formally introduce a few key notions and notations of Goyal and Vega-Redondo (2007). A strategy for player i is a set s_i of agents with whom i would like to link. Hence, $j \in s_i$ reads as agent i intends to form a link with j, and $j \notin s_i$ reads as agent i does not intend to form a link with j. The set of all pure strategies of agent i is S_i. A strategy profile $s = (s_1, \ldots, s_n)$ maps to an (undirected) network $g(s)$ in the following way. A link $i - j$ exists in some network if, and only if, the link was mutually intended, $i \in s_j$ and $j \in s_i$. The agents with whom i has links in the network $g(s)$ is referred to as $g_i(s)$.

In any network, each unordered pair of connected agents generates a unit surplus. This surplus is then divided equally among the agents essential to the connection. Specifically, an agent k is *crucial for connecting i and j* if all paths between i and j pass through k. For any unordered pair (i, j), $E(i, j; g)$ is the set of agents who are crucial for connecting i and j in the network g; let $e(i, j; g) = |E(i, j; g)|$ be the cardinality of that set. In Goyal and Vega-Redondo (2007), individual payoffs are given by:

$$u_i(s) = \sum_{j \in N \setminus \{i\}} \frac{1}{e(i, j; g(s)) + 2} + \sum_{(j,k) \in (N \setminus \{i\}) \times (N \setminus \{i\})} \frac{\gtrdot 1_{i \in E(j,k;g(s))}}{e(j, k; g(s)) + 2} - c|g_i(s)|, \tag{3}$$

where $c > 0$ is the marginal cost of a link. Let me contrast the above payoff with the following linear specification of the current payoff in (1):

$$u_i(s) = (\Delta_{ii}(g(s)) - 1) + \sum_{j \in N \setminus \{i\}} \frac{(\Delta_{ij}(g(s)) - 1)}{2} - c|g_i(s)|. \tag{4}$$

so that an agent earns 1 every time they are essential in a connection.

The equilibrium concept is the *strict bilateral equilibrium* (SBE). A strategy profile s is an SBE if, and only if,

- For each $i \in N$, for every strategy $s_i' \in S_i$ such that $g(s_i', s_{-i}) \neq g(s)$, it holds that $u_i(s) > u_i(s_i', s_{-i})$.
- For any pair of agents $i, j \in N$ and every strategy pair $(s_i', s_j') \in S_i \times S_j$ such that $g(s_i', s_j', s_{-i-j}) \neq g(s)$, $u_i(s_i', s_j', s_{-i-j}) \geq u_i(s)$ implies $u_j(s_i', s_j', s_{-i-j}) < u_j(s)$.

Despite the different nature of benefits, equilibrium predictions converge when the number of players is sufficiently large: an SBE is either a star network or the empty network. The reasons, however, differ from those in Goyal and Vega-Redondo (2007). With rival benefits, there is an additional incentive to form links: to reduce dependence on intermediaries. However, when the set of agents is sufficiently large, a significant reduction of dependence is too costly, so that the players coordinate on the centralized architecture of a star. The following observations apply to the current model and explain why the star network remains an SBE when intermediation benefits are non-rival.

1. First, an SBE is *acyclic*. Indeed, the value of a link in a cycle is to decrease the number of other crucial agents. However, the non-rival nature of benefits renders such links worthless to the agents who maintain them. Therefore, these agents would strictly improve their payoffs by cutting them.
2. Second, a non-empty SBE is *connected*. If the network has two acyclic components A and B that each encompass several individuals, two agents with a link to a leaf in their respective components would do better by cutting those links and then forming one link between them. The deviation keeps the number of links paid for by the two agents constant; however, it yields the same or higher benefits for both. Next, consider a network with one acyclic component and some isolated singletons. Note that an agent i in the acyclic component who maintains a link with a leaf j has a payoff-equivalent deviation, which is first to remove the link with j and then add a link with k. If k accepts the link proposal, the latter taps the same payoff as the leaf j in the original network. Hence, if the leaf j prefers not to cut the link with i in the original network, k should accept i's link proposal. It follows that the original network is not an SBE.
3. Last, *a non-empty SBE is a star network* (see Appendix H for a formal proof). To illustrate the argument concretely, consider the network on the left in Fig. 11. If the network is an SBE, the leaf agent 4 must prefer not to cut the link with agent 1 – i.e., $c < 4$. Let agent 3, who is relatively more central than agent 4, first cut their link with agent 1 and offer simultaneously to agent 4 to form a link. If agent 4 accepts, the network will be as shown on the right in Fig. 11. Because agent 1 remains reachable from agent 3, the deviation is payoff-equivalent to the latter. Accepting the link would increase agent 4's centrality. Its marginal benefit is 4 – as agent 4 is now crucial for connecting agents 2 and 1 to agents 3 and 5. Therefore, the network is an SBE only if agent 4 prefers not to accept agent 3's link proposal, or $c > 4$. Recall that for such values of the cost, agent 4 prefers to cut their link with agent 1. Thus, the network is not an SBE.

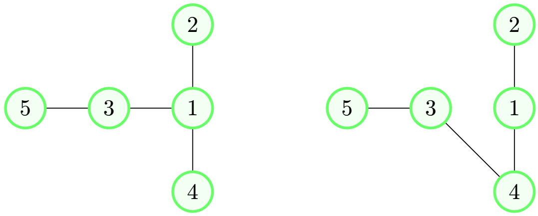

Fig. 11 Profitable deviations in this game

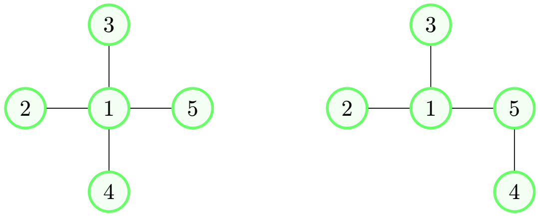

Fig. 12 A star network on the left; on the right, a deviation by agents 4 and 5

The conditions under which a star network is an SBE markedly differ across the two models. In Goyal and Vega-Redondo (2007), the most profitable deviations exploit the additional incentive created by rival benefits, which is to form a link between peripheral agents in order to decrease their dependence on the central agent. As argued in point 1 above, this incentive is absent with non-rival benefits and therefore does not threaten the existence of stars as equilibria. However, as explained in point 3 above, two peripheral agents in a star can team up to improve the centrality of one of them and the associated intermediation benefits, as illustrated in Fig. 12. Whether the star is an SBE, hence, depends on the profitability of these deviations.

The following Corollary summarizes the difference in results between this game and Goyal and Vega-Redondo (2007)'s.

Corollary 11 *(See Theorem 1 pp. 468 in Goyal and Vega-Redondo (2007).) Consider the payoffs in* (3) *and* (4)*, and suppose that* n *is sufficiently large. (1) An SBE is either a star network or the empty network. (2) In Goyal and Vega-Redondo (2007)'s game, if* $c > \frac{1}{6}$*, a star network is the unique non-empty SBE. The empty network is an SBE if, and only if,* $c > \frac{1}{2}$*. (3) In this game, if* $c > n - 2$*, the star is the unique non-empty SBE. The empty network is an SBE if, and only if,* $c > 1$.

5 Conclusion

This paper studies a non-cooperative game of network formation in which agents value not only their own access to the network, but also the contribution their links make to others' access. Link formation, therefore, reflects a trade-off between the social value of links and the private cost of forming them. I characterize the architecture of strict Nash networks and the conditions under which those architectures arise in equilibrium.

Under directed links, strict Nash networks may be either flat or hierarchical. Flat equilibria consist of a wheel that may or may not encompass all agents, whereas hierarchical equilibria are trees in which some agents act as intermediaries between otherwise disconnected groups. For sufficiently low link-formation costs, exhaustive wheels are the unique equilibrium architecture, and best-response dynamics converge to them with probability one. For intermediate costs, multiple strict Nash architectures may coexist, although exhaustive wheels and trees Pareto-dominate some alternative equilibrium architectures.

These findings contrast with the seminal model of Bala and Goyal (2000), where agents value their links solely for the private benefit of reaching others. Because intermediation does not generate benefits in their model, a non-empty equilibrium network must be connected. In the present game, by contrast, agents may benefit from occupying positions that facilitate access for others. As a result, equilibrium networks may be disconnected, and non-empty networks arise for a wider range of link-formation costs.

The comparison with Goyal and Vega-Redondo (2007) highlights a different mechanism. In their model, intermediation benefits are rival: agents have incentives to form links that bypass existing intermediaries, thereby reducing others' intermediation benefits. In the present game, these benefits are non-rival, so there are no such incentives. Nevertheless, when links are undirected, both models select star architectures in equilibrium. However, the conditions under which the star is an equilibrium differ across the two environments. In Goyal and Vega-Redondo (2007), equilibrium requires peripheral agents not to gain from forming a link that bypasses the center. Here, equilibrium requires that peripheral agents cannot profitably rearrange links to increase their own centrality and the associated intermediation benefits.

Finally, the paper's predictions are robust to alternative specifications of non-rival intermediation benefits. In particular, the same equilibrium architectures arise when network benefits are defined purely in terms of intermediation, so that agents derive no direct benefit from accessing or being accessible to others. This suggests that the architecture of equilibrium networks is driven less by the precise specification of non-rival benefits than by the broader distinction between social and private incentives in link formation.

Acknowledgements I am grateful to the editor and two anonymous referees for their constructive comments and suggestions, which have significantly improved the paper. I thank Francis Bloch, Sidartha Gordon, Nicolas Klein, László Kóczy, Ming Li, Szilvia Pápai, and Françoise Forges for helpful comments.

Funding Open access funding provided by Corvinus University of Budapest.

Appendix A diffusion auction

The reader may have noticed that the mechanism presented in Section 2.2 is not generally budget balanced. To address this issue, suppose that the auctioneer sets a reserve price, $\gamma \in [0, 1)$. The mechanism remains incentive compatible and individually rational. The mechanism remains incentive compatible and individually rational. Hence, agents report the network truthfully and every informed agent participates. Therefore, given an information seed $j_0 \in N$, the set of participants coincides with the set of agents whom j_0 can reach, denoted $R(j_0; g)$.

The purpose of this appendix is threefold. First, I derive the interim expected payoff generated by the mechanism with reserve price γ. Second, I verify that the resulting payoff satisfies Assumption 1. Third, I show that an appropriate reserve price guarantees budget balance when the number of agents is finite.

To derive the expected payoff generated by the mechanism, fix an information seed $j_0 \in N$ and a network g. All variables defined below are conditional on this choice of j_0 and g. For any $i \in R(j_0; g)$, agent i's utility from the mechanism is:

$$\pi_i(t) = v^*_{(g)} - \max\{\gamma, v^*_{(\tilde{g}^i)}\},$$

where

$$v^*_{(g)} := \max_{j \in R(j_0; g)} v_j$$

and

$$v^*_{(\tilde{g}^i)} := \max_{j \in R(j_0; \tilde{g}^i)} v_j.$$

Interim Expected Utility from the Mechanism with Reserve Price γ

To derive the expected payoff generated by the mechanism, fix an information seed $j_0 \in N$ and a network g. It is useful to partition the agents reachable from j_0 into two sets: those who remain reachable when agent i is removed and those whose participation depends on agent i. Agent i's transfer depends on the highest valuations in these two disjoint sets.

For each $i \in R(j_0; g)$, define:

- The highest valuation among the agents who can be reached from j_0 only through agent i,
$$V(i) = \max_{j \in R(j_0; g) \setminus R(j_0; \tilde{g}^i)} v_j.$$
- The highest valuation among the agents who remain reachable when agent i and their links are removed,
$$V(-i) = \max_{j \in R(j_0; \tilde{g}^i)} v_j.$$
- Let
$$U^+ := \{V(i) \geq V(-i) \text{ and } V(i) > \gamma\}.$$

  Agent i receives a positive transfer only if the winner belongs to the set of agents whose participation depends on i and the winning valuation exceeds the reserve price. The event U^+ captures precisely these circumstances.

To simplify notation, I suppress the dependence of r_{j_0} and Δ_{ij_0} on the network g. Recall that

$$\Delta_{ij_0} = |R(j_0; g) \setminus R(j_0; \tilde{g}^i)|$$

and

$$r_{j_0} - \Delta_{ij_0} = |R(j_0; \tilde{g}^i)|.$$

The probability of the event U^+ is

$$\mathbb{P}(U^+) = \frac{\Delta_{ij_0}}{r_{j_0}}\big(1 - \gamma^{r_{j_0}}\big).$$

The first term is the probability that the highest valuation among the reachable agents belongs to an agent whose participation depends on i. The second term is the probability that this valuation exceeds the reserve price.

An expression for the interim expected utility of agent i is therefore:

$$\sum_{j_0 \in N} \frac{1}{n} \times \mathbb{P}(U^+)\Bigg(\mathbb{E}[V(i)|U^+] - \mathbb{P}(V(-i) \leq \gamma|U^+) \times \gamma$$

$$- \mathbb{P}(V(-i) > \gamma|U^+) \times \mathbb{E}[V(-i)|U^+, V(-i) > \gamma]\Bigg).$$

The first term gives the expected welfare generated when the winner belongs to the set of agents whose participation depends on i. The second term corresponds to the case in which all agents who remain reachable without i have valuations below the reserve price. In that case, the welfare would have been γ. The last term corresponds to the case in which some agent remains reachable without i and has a valuation exceeding the reserve price. The transfer to i then depends on the expected value of the highest such valuation.

The derivation proceeds by evaluating the three objects that appear in the expression above.

1. **Computing** $\mathbb{E}[V(i) \mid U^+]$**.**
   Conditional on U^+, the winner is the agent with the highest valuation among the r_{j_0} reachable agents and this valuation exceeds the reserve price. Hence $\mathbb{E}[V(i) \mid U^+]$ coincides with the expected maximum valuation among r_{j_0} independent draws from the uniform distribution conditional on exceeding γ:

$$\mathbb{E}[V(i)|U^+] = \mathbb{E}\left[\max_{j \in R(j_0;g)} v_j \;\middle|\; \max_{j \in R(j_0;g)} v_j > \gamma\right].$$

   The conditional density of the maximum is

$$f_{(1)|\gamma}(x) = \frac{r_{j_0} x^{r_{j_0}-1}}{\int_\gamma^1 r_{j_0} x^{r_{j_0}-1} dx} = \frac{r_{j_0} x^{r_{j_0}-1}}{1 - \gamma^{r_{j_0}}}.$$

   Therefore,

$$\begin{aligned}\mathbb{E}[V(i)|U^+] &= \int_\gamma^1 x f_{(1)|\gamma}(x)dx \\ &= \frac{r_{j_0}}{1-\gamma^{r_{j_0}}}\int_\gamma^1 x^{r_{j_0}}dx \\ &= \left(\frac{r_{j_0}}{r_{j_0}+1}\right)\left(\frac{1-\gamma^{r_{j_0}+1}}{1-\gamma^{r_{j_0}}}\right).\end{aligned}$$

2. **Computing** $\mathbb{P}(V(-i) \le \gamma|U^+)$**.**
   This is the probability that all agents who remain reachable without i have valuations below the reserve price, conditional on U^+:

$$\mathbb{P}(V(-i) \le \gamma|U^+) = \frac{\mathbb{P}(V(-i) \le \gamma, U^+)}{\mathbb{P}(U^+)} = \frac{\mathbb{P}(V(-i) \le \gamma < V(i))}{\mathbb{P}(U^+)}.$$

   Since there are $(r_{j_0} - \Delta_{ij_0})$ agents reachable without i and Δ_{ij_0} agents whose participation depends on i,

$$\mathbb{P}(V(-i) \le \gamma|U^+) = \frac{\gamma^{r_{j_0}-\Delta_{ij_0}}(1-\gamma^{\Delta_{ij_0}})}{\mathbb{P}(U^+)}.$$

3. **Computing** $\mathbb{P}(V(-i) > \gamma|U^+)\mathbb{E}[V(-i)|U^+, V(-i) > \gamma]$**.**
   This term corresponds to the case in which some agent remains reachable without i and has a valuation exceeding the reserve price.

$$\begin{aligned}&\mathbb{P}(V(-i) > \gamma|U^+) \times \mathbb{E}[V(-i)|U^+, V(-i) > \gamma] \\ &= \frac{\mathbb{P}(V(i) \ge V(-i) > \gamma)}{\mathbb{P}(U^+)} \times \int_\gamma^1 x\, d\left(\frac{\mathbb{P}(V(-i) \le x, V(i) \ge V(-i) > \gamma)}{\mathbb{P}(V(i) \ge V(-i) > \gamma)}\right) \\ &= \frac{1}{\mathbb{P}(U^+)}\int_\gamma^1 x(r_{j_0} - \Delta_{ij_0})x^{r_{j_0}-\Delta_{ij_0}-1}(1-x^{\Delta_{ij_0}})dx \\ &= \frac{1}{\mathbb{P}(U^+)}\left[\left(\frac{r_{j_0}-\Delta_{ij_0}}{r_{j_0}-\Delta_{ij_0}+1}\right)(1-\gamma^{r_{j_0}-\Delta_{ij_0}+1}) - \left(\frac{r_{j_0}-\Delta_{ij_0}}{r_{j_0}+1}\right)(1-\gamma^{r_{j_0}+1})\right].\end{aligned}$$

Replacing the three expressions above into the interim utility yields:

$$\begin{aligned}\frac{1}{n}\sum_{j_0\in N}\Bigg(&\frac{\Delta_{ij_0}}{(r_{j_0}+1)}(1-\gamma^{r_{j_0}+1}) - (\gamma^{r_{j_0}-\Delta_{ij_0}+1} - \gamma^{r_{j_0}+1}) \\ &- \left(\frac{r_{j_0}-\Delta_{ij_0}}{r_{j_0}-\Delta_{ij_0}+1}\right)(1-\gamma^{r_{j_0}-\Delta_{ij_0}+1}) + \left(\frac{r_{j_0}-\Delta_{ij_0}}{r_{j_0}+1}\right)(1-\gamma^{r_{j_0}+1})\Bigg)\end{aligned}$$

The above can be re-arranged as:

$$\frac{1}{n}\sum_{j_0\in N}\left(\left(\frac{r_{j_0}+\gamma^{r_{j_0}+1}}{r_{j_0}+1}\right)-\left(\frac{r_{j_0}-\Delta_{ij_0}+\gamma^{r_{j_0}-\Delta_{ij_0}+1}}{r_{j_0}-\Delta_{ij_0}+1}\right)\right)$$

This expression depends on the reach of the information seed and on agent i's contribution to it, as in the reduced-form payoff introduced in Section 2.1.

Assumption 1

To verify Assumption 1, it is sufficient to show that $h(z)$ is strictly increasing, where $h(z)=\frac{z+\gamma^{z+1}}{z+1}$ for $z\in\mathbb{Z}_+$.

$$h(z+1)-h(z)=\frac{1+\gamma^{z+2}(z+1)-\gamma^{z+1}(z+2)}{(z+1)(z+2)}>\Big|_{\gamma=1}0.$$

The strict inequality owes to the fact that the ratio is decreasing in the value of γ and $\gamma\in[0,1)$.

Budget Balance

I show that introducing a sufficiently high reserve price guarantees budget balance when the number of agents is finite. The argument proceeds by deriving a lower bound on the auctioneer's net revenue that holds for every realization of valuations. This lower bound depends only on the reserve price and the number of agents. It follows that, when the reserve price is sufficiently high, the auctioneer's net revenue is non-negative regardless of the network structure or the realized valuations.

Fix an information seed j_0 and a network g. If no participant reports a valuation above the reserve price γ, the item is not sold and the auctioneer's net revenue is equal to zero. Suppose instead that the item is sold, and let w denote the winner.

By construction of the mechanism, the winner pays at least γ to the auctioneer. Consider now the transfers paid to the remaining participants. For any participant $i\neq w$, the transfer received by i is equal to the welfare loss that would result from removing i and their links. Since

$$\pi_i=v_w-\max\{\gamma,v^*_{(\tilde{g}^i)}\}\quad\text{and}\ \max\{\gamma,v^*_{(\tilde{g}^i)}\}\geq\gamma,$$

it follows that

$$\pi_i\leq v_w-\gamma\leq 1-\gamma.$$

Since there are at most $n-1$ participants other than the winner, the auctioneer's net revenue is bounded below by

$$\gamma-(n-1)(1-\gamma).$$

This lower bound is non-negative whenever

$$\gamma\geq\frac{n-1}{n}.$$

Since the number of agents is finite, $(n-1)/n < 1$. Therefore, there exists a reserve price $\gamma \in [0, 1)$ such that the auctioneer's net revenue is non-negative for every realization of valuations.

Consequently, for any finite number of agents, a sufficiently high reserve price guarantees budget balance.

Appendix B Proof of Corollary 2

STATEMENT 1: Consider a network g with a component A that is not a maximal element of $\hookrightarrow$ and is not a singleton. Because A is not maximal, there exists B, another component in g, such that $B \hookrightarrow A$. Since A is not a singleton, $|A| \geq 2$. Since $B \hookrightarrow A$ and A is not a singleton, there exist $i \in B$ and distinct agents $j, k \in A$ such that $j \in g_i$ and $j \in g_k$. Since i and k do not belong to the same component, we conclude from Proposition 1 that g is not Nash.
STATEMENT 2. Suppose that some network g admits a component A such that $|A| > 1$ and A is not minimally connected. Then there exist distinct agents $i, k \in A$ such that $k \in g_i$ and A remains a component in the network $g' = (g'_i, g_{-i})$, where $g'_i = g_i \setminus \{k\}$. Because A is a component and $|A| > 1$, there exists an agent $j \in A$, distinct from i and k, such that $k \in g_j$. Moreover, since deleting $i \to k$ does not disrupt the component A, agent i still has a path to j in g'; equivalently, in g, there is a path from i to j that does not pass through k. Yet $k \in g_i$ and $k \in g_j$. This contradicts Proposition 1. Therefore, g is not Nash.
STATEMENT 3. Let A and B be two distinct maximal components in some network g. The proof is by contradiction. Suppose that g is Nash and neither A nor B is an isolated singleton.

Take any agents $a \in A$ and $b \in B$. Since A and B are maximal, a deviation by an agent in A (respectively, B) can only affect the reach of agents in A (respectively, B). Moreover, Proposition 1 implies that A and B cannot access a common component. Therefore, the sets of agents reachable from a and b are disjoint. Denote these sets by $N_a(g)$ and $N_b(g)$, respectively. Without loss of generality, suppose that $r_a(g) \geq r_b(g)$.

Consider the following deviation by agent b. Let $g'_b = \{a\}$ and let $g' = (g'_b, g_{-b})$. The deviation replaces b's existing links with a single link to agent a. Since B is not an isolated singleton, $|g_b| \geq 1$, and therefore

$$|g'_b| \leq |g_b|.$$

It remains to show that agent b's network benefit strictly increases.

- For agent b,
$$r_b(g') = r_a(g) + 1 > r_b(g).$$

  Indeed, agent b can access all agents reachable from a and agent b. Since $r_b(.) = \Delta_{bb}(.)$ in every network, it follows from Assumption 1 that

$$\Phi(r_b(g'), \Delta_{bb}(g')) > \Phi(r_b(g), \Delta_{bb}(g)).$$

The proof is complete if $B = \{b\}$. Suppose instead that B contains another agent.

- Consider any $i \in B \setminus \{b\}$. Since i can still reach b in the network g', it holds that

$$r_i(g) = r_i(\tilde{g}^b) + \Delta_{bi}(g)$$

and

$$r_i(g') = r_i(\tilde{g}^b) + \Delta_{bi}(g').$$

Hence,

$$\Phi\big(r_i(g'), \Delta_{bi}(g')\big) = \Phi\big(r_i(g) + \big(\Delta_{bi}(g') - \Delta_{bi}(g)\big),\ \Delta_{bi}(g) + \big(\Delta_{bi}(g') - \Delta_{bi}(g)\big)\big).$$

By Assumption 1, this expression is strictly greater than

$$\Phi(r_i(g), \Delta_{bi}(g))$$

whenever

$$\Delta_{bi}(g') - \Delta_{bi}(g) > 0.$$

To verify this inequality, first note that

$$\Delta_{bi}(g') \geq r_a(g) + 1,$$

because in g' agent b is essential to agent i for reaching both agent b and all agents in $N_a(g)$.
Second, since i and b belong to the same component in g,

$$r_i(g) = r_b(g).$$

Moreover, agent b is never essential to agent i for reaching agent i when $i \neq b$. Hence,

$$\Delta_{bi}(g) \leq r_b(g) - 1.$$

Combining these inequalities with the assumption that $r_a(g) \geq r_b(g)$ yields

$$\Delta_{bi}(g') - \Delta_{bi}(g) \geq r_a(g) + 1 - r_b(g) + 1 = r_a(g) - r_b(g) + 2 > 0.$$

Therefore,

$$\Phi(r_i(g'), \Delta_{bi}(g')) > \Phi(r_i(g), \Delta_{bi}(g))$$

for every $i \in B \setminus \{b\}$.

Consequently, agent b's deviation is strictly profitable. This contradicts that g is Nash.

Appendix C Proof of Proposition 3

Observation 1 *In a strict Nash network, a component with several agents forms a wheel.*

By Corollary 2 Statement 2, a component A with several agents is minimally connected if the network g is Nash. If A does not form a wheel, $\exists i, j, k \in A$ such that $i \neq j \neq k$ and $k \in g_i, k \in g_j$. Note, agent i has a payoff-equivalent strategy: $g'_i = (g_i \setminus \{k\}) \cup \{j\}$. Thus g is not strict Nash.

Observation 2 *If g is strict Nash and there exists a component A that is maximal and not minimal, then (a) $|A| > 1$ and (b) A is the greatest element of $\hookrightarrow$ (i.e., $A \hookrightarrow B$ for all components B).*

Statement (a): *Let g be a network with a component $A = \{a\}$ that is maximal and not minimal. Thus, $|g_a| \geq 1$. Also, there exists a component B that is minimal and not maximal (in particular, $A \hookrightarrow B$). If g is Nash, $B = \{b\}$ (by Corollary 2 Statement 1). Thus, $|g_b| = 0$. Because A is maximal, not minimal and $A = \{a\}$, $u_a(g) = \Phi(k, k) - C(|g_a|)$ for some $k \in \{2, \ldots, n\}$. If g is strict Nash, a must incur a loss if they cut all their links: $u_a(g) = \Phi(k, k) - C(|g_a|) > \Phi(1, 1) = u_a(g'_a = \emptyset, g_{-a})$. Because the cost function C is increasing and $|g_a| \geq 1$, $C(|g_a|) \geq C(1)$. Thus, a necessary condition for g to be strict Nash is:*

$$C(1) < \Phi(k, k) - \Phi(1, 1). \tag{C1}$$

Consider now agent b; let them deviate to $g'_b = \{a\}$, and let g' be the resulting network. Note, if $r_i(g') \neq r_i(g)$ for some $i \in N$, then i can reach b in g and $r_i(g') \geq r_i(g)$ and the inequality is strict for at least one agent (namely, b). Also, $\Delta_{bi}(g') - \Delta_{bi}(g) = r_i(g') - r_i(g)$ for all $i \in N$. Assumption 1 part 3 allows us to conclude that $u_b(g') - u_b(g) \geq \Phi(k, k) - \Phi(1, 1) - C(1)$. If g is strict Nash, b's deviation must decrease their payoff. This is the case only if:

$$C(1) > \Phi(k, k) - \Phi(1, 1). \tag{C2}$$

The two necessary conditions in (C1) and (C2) cannot hold simultaneously. Hence, g is not strict Nash.

Statement (b): *Suppose some network g has some, but not all, components comparable via $\hookrightarrow$. Then, there are two components B and B' such that B is minimal, not maximal; B' is maximal and minimal. If g is Nash, B is a singleton and B' is an isolated singleton (by Corollary 2 Statements 1 and 3). Set $B = \{b\}$ and $B' = \{b'\}$. By Proposition 1, there is a unique link that points towards b; let i be the agent who maintains it. Clearly, agent i has a payoff-equivalent strategy, $g'_i = (g_i \setminus \{b\}) \cup \{b'\}$. Thus, g is not strict Nash.*

Observation 3 *In a strict Nash network in which no two components are comparable via $\hookrightarrow$, a wheel component encompasses at least three agents.*

Consider a network with a wheel component, and that has no two components comparable via $\hookrightarrow$. If g is Nash, each component other than the wheel component is *an isolated singleton (by Corollary 2 Statement 3). Suppose there are only two agents in the wheel, i and j. The payoff to either agent is:*

$$\Phi(2,2)+\Phi(2,1)-C(1).$$

Let agent i deviate as follows: i cuts his link with j and simultaneously adds a link with some isolated singleton k. The resulting network has a chain, $j \to i \to k$. i's payoff after deviating is:

$$\Phi(3,2)+\Phi(2,2)-C(1).$$

The deviation is strictly profitable since, by Assumption 1 part 3, $\Phi(3,2) > \Phi(2,1)$. Hence g is not strict Nash.

Proposition 3 follows directly from these three observations:

- *If g is strict Nash and g has one component, g is an exhaustive wheel (by Observation 1).*
- *If g is strict Nash, no two components of g are comparable via $\hookrightarrow$, g is either the empty network or a non-exhaustive wheel (by Observations 1 and 3, by Corollary 2 Statement 3).*
- *If g is strict Nash and there are some distinct components comparable via $\hookrightarrow$, then g is either a root-leaves tree or a tree with intermediaries (by Observations 1 and 2, by Corollary 2).*

Appendix D Proof of Proposition 4

For each architecture, the proof identifies the most profitable deviations. Any deviation not listed in Proposition 4 is payoff inferior to one of the deviations considered below and therefore cannot be a best response.

The statement for the empty network can be directly verified. Let g be an exhaustive wheel. In Proposition 4 Statement (2), the deviations g'_i such that $|g'_i| \geq |g_i|$ have not been considered (for any $i \in N$). This is because they are never best responses. To see why, consider a line network defined by the sequence of links $1 \to 2 \ldots \to n$. Let agent n play some strategy g'_n that prescribes forming a link with some agent $i \neq 1$. In the resulting network g', there are two links towards i (the one maintained by $i-1$ and n's). Also, either $i-1$ and n are in different components or there is a path from n to $i-1$ that does not go through i. From there, similar steps as in the proof of Proposition 1 can be followed to show that g'_n is not a best response for n.

Consider a non-exhaustive wheel g with $w \in \{3,\ldots,n-1\}$ agents. Throughout, A refers to the wheel component. Let i be an isolated singleton; let them deviate and form links with some set of agents g'_i. Deviations such that:

- *$g'_i \cap A = \emptyset$ are never best responses, as they yield a lower payoff than the deviations that have one link with a wheel member and $|g'_i|-1$ links towards isolated agents (same cost of link formation, strictly higher benefit because $w > 1$).*
- *$|g'_i \cap A| > 1$ are never best responses, as they yield a lower payoff than the deviations that have one single link towards the wheel and the same number of links with isolated agents (lower cost of link formation, same benefit).*

It is for this reason that the deviations presented in Proposition 4 statement (3.a) are of the form $g'_i = \{a\} \cup L$, *where* $a \in A$ *and* $L \cap A = \emptyset$. *I.e., there is one link towards the wheel; if* $|g'_i| > 1$, *the rest of the links are all towards isolated agents.*

Next, let i *be a wheel member. The deviations (denoted by* g'_i*) that are not considered in Proposition 4 are all such that:* $\exists a \in g'_i$ *such that* $a \notin g_i$, $a \in A$. *Let* g' *be the resulting network. In* g', *there are two links pointing towards* a*: one maintained by* i, *the other by the player* $j \in A$ *for whom* $g_j = \{a\}$. *Note that in* g', *either* i *and* j *are in two different components, or there is a path from* i *to* j *that does not go through* a. *From there, similar steps as in the proof of Proposition 1 can be followed to show that* g'_i *is not a best response for* i.

Let g *be a root-leaves tree.* A *denotes the root (wheel) component; consider any* $i \in A$. *The deviations of* i *(denoted by* g'_i*) that are not considered in Proposition 4 Statement (4.b) satisfy one of the following:*

- $|g'_i| > |g_i|$.
- $|g'_i| \leq |g_i|$, $|g_i| \geq 1$ *and* $g'_i \nsubseteq g_i$.

Let g' *be the network defined by the strategy profile* (g'_i, g_{-i}), *assuming that* g'_i *satisfies one of the above properties. In* g', $\exists j \in g'_i$ *such that* j *has two links pointing towards them and one is maintained by* i, *the other by the agent* k *for whom* $j \in g_k$; *and the conclusion of Proposition 1 is violated. As such, one can use the proof of Proposition 1 to reach the desired conclusion.*

Last, consider a leaf i *in* g. *A deviation* g'_i *consists in forming links. Observe that only the agents in the root can access* i, *and* $r_j(g) = n$ *for each* $j \in A$. *Hence, for* g' *the network that results when leaf* i *deviates to some* g'_i, $\Delta_{ij}(g') - \Delta_{ij}(g) = r_j(g') - r_j(g) = 0$ *for all* $j \neq i$. *The maximum gain in payoff that* i *can earn by deviating is then* $\Phi(n, n) - C(1) - \Phi(1, 1)$*; this is the gain to* i *from forming one link with any agent in* A.

Appendix E Proof of Corollary 5

Statement 1 follows from the following observations.

Observation 1 *If a non-exhaustive wheel is strict Nash, so are the exhaustive wheels.*

For $w \in \{3, \ldots, n\}$, *consider a flat network* g^w *that has one wheel component with* w *agents and the rest of the components are isolated singletons. An exhaustive wheel has* $w = n$*; a non-exhaustive wheel has* $w \in \{3, \ldots, n-1\}$. *For any given* w *in the appropriate range,* g^w *is strict Nash only if no agent in the wheel component wants to cut their link. Mathematically,*

$$C(1) < \sum_{k=1}^{w-1} \left(\Phi(w, w-k+1) - \Phi(k, 1)\right).$$

When $w = n$, the above is both necessary and sufficient. Therefore, it suffices to show that the right side of the above inequality is increasing in w. Indeed,

$$\sum_{k=1}^{w} (\Phi(w+1, w-k+2) - \Phi(k,1)) - \sum_{k=1}^{w-1} (\Phi(w, w-k+1) - \Phi(k,1))$$
$$= \Phi(w+1,2) - \Phi(w,1) + \sum_{k=1}^{w-1} (\Phi(w+1, w-k+2) - \Phi(w, w-k+1))$$

Each term in this sum is positive by Assumption 1 part 3. This finishes the proof.

Observation 2 *If a non-exhaustive wheel is strict Nash, so is the empty network.*
The empty network is strict Nash if, and only if, for all $x \in \{1, \ldots, n-1\}$,

$$\Phi(1,1) > \Phi(x+1, x+1) - C(x). \tag{E3}$$

A non-exhaustive wheel is strict Nash only if for all $x \in \{1, \ldots, n-w\}$,

$$\Phi(1,1) > \Phi(w+x, w+x) - C(x). \tag{E4}$$

For all $x \in \{1, \ldots, n-w\}$, the right side of (E3) is strictly lower than the right side of (E4), since $w > 1$. For $x \in \{n-w+1, \ldots, n-1\}$, an isolated singleton in a non-exhaustive wheel cannot form links to more than the $n-w$ agents outside the wheel once they have formed a link to the wheel. The most profitable deviation is therefore the one in which the isolated singleton links to the wheel and to all other isolated singletons, yielding payoff $\Phi(n,n) - C(n-w)$. By Assumption 1 part 3 and the monotonicity of C, this payoff is strictly higher than $\Phi(x+1, x+1) - C(x)$. Therefore, if (E4) holds, so does (E3).

Observation 3 *If a tree (root-leaves or with intermediaries) is strict Nash, so is an exhaustive wheel.*

Consider some strict Nash tree g. In g, let A denote the root component, and let $B = \{b \in N : \ g_b \neq \emptyset, b \notin A\}$ be the set of intermediaries. Recall, A forms a wheel.

Construct an exhaustive wheel g' as follows. Start from the wheel formed by the root A. For each agent $a \in A$ with links with agents in B, replace the link from a to the next agent $a' \in A$ by a path that visits all branches originating from a. More precisely, if the branches originating from a are

$$a \to a_1^1 \to \cdots \to a_{m_1}^1, \quad \ldots, \quad a \to a_1^\ell \to \cdots \to a_{m_\ell}^\ell,$$

then replace the link from a to the next wheel member a' by

$$a \to a_1^1 \to \cdots \to a_{m_1}^1 \to a_1^2 \to \cdots \to a_{m_2}^2 \to \cdots \to a_1^\ell \to \cdots \to a_{m_\ell}^\ell \to a'.$$

Repeating this operation for every root member in A yields an exhaustive wheel. By construction, every agent who is essential for some connection in g remains essential for that connection in g'.

If the tree g is strict Nash, no $i \in A \cup B$ wants to cut all their links. Notice that (1) in the tree g, $|g_i| \geq 1$ for each $i \in A \cup B$; furthermore, (2) the set of connections for which i is essential in g is a subset of the set of connections for which i is essential in g', for all $i \in A \cup B$. Therefore, for each $i \in A \cup B$,

$$u_i(g) - u_i(g_i = \emptyset, g_{-i}) \leq u_i(g') - u_i(g_i = \emptyset, g'_{-i}).$$

The tree g is strict Nash only if the expression on the left hand side of the above inequality is strictly positive: no agent with links in the tree must gain from cutting all their links. If such is the case, $u_i(g') - u_i(g_i = \emptyset, g'_{-i}) > 0$ which is equivalent to g' being strict Nash (by Proposition 4 statement 2).

Observation 4 *If a tree (root-leaves or with intermediaries) is strict Nash, so is the empty network.*

A tree is strict Nash only if no leaf wants to form one link towards the root component of their network. A minimum on the payoff difference from such a deviation is $\Phi(n, n) - \Phi(1, 1) - C(1)$. Thus, a necessary condition for a tree to be strict Nash is that $\Phi(n, n) - C(1) < \Phi(1, 1)$. Assume this holds; we will show that no deviation is then profitable in the empty network. Indeed, if any $i \in N$ forms links in the empty network, their payoff is: $\Phi(x + 1, x + 1) - C(x)$ for some $x \in \{1, \ldots, n - 1\}$. Because C is strictly increasing, $\Phi(x + 1, x + 1) - C(x) \leq \Phi(x + 1, x + 1) - C(1)$. Assumption 1 part 3 establishes that $\Phi(x + 1, x + 1) - C(1) \leq \Phi(n, n) - C(1)$. Given the necessary condition on $C(1)$ tied to the trees, $\Phi(n, n) - C(1) < \Phi(1, 1)$. Gathering, we have $\Phi(x + 1, x + 1) - C(x) < \Phi(1, 1)$ for any $x \in \{1, \ldots, n - 1\}$ (recall that $\Phi(1, 1)$ is the payoff of any agent in the empty network). It follows that the empty network is strict Nash.

Statement 2: It can be verified that the benefit an agent taps in a wheel component increases with its size. The remainder shows that the payoff to an isolated singleton, $\Phi(1, 1)$, is strictly less than the payoff in an exhaustive wheel when it is strict Nash. Indeed, an exhaustive is strict Nash if, and only if,

$$\sum_{k=1}^{n} \Phi(n, k) - C(1) > \sum_{k=1}^{n} \Phi(k, 1).$$

The result follows from the observation that $n > 1$ and by Assumption 1 part 1, $\Phi(k, 1) > 0$ for each k from 2 to n.

Statement 3. In a strict Nash tree, every agent with links strictly prefers their equilibrium strategy to cutting all links. The payoff from cutting all links is at least $\Phi(1, 1)$, which implies that each agent with links earns strictly more than in the empty network. Moreover, each leaf earns more than $\Phi(1, 1)$, since they further benefit from being accessible to other agents. Hence every agent is strictly better off in the tree than in the empty network. Therefore, the tree payoff-dominates the empty network.

Appendix F Proof of Proposition 6

The proof proceeds in two steps. First, I assume that some agent reaches the entire population, and I show that there exists a sequence of best responses that leads to an exhaustive wheel. Second, I show that this assumption is without loss of generality.

Consider any initial network $g(0)$. Assume first that $r_i(g(0)) = n$ for some $i \in N$. The proof proceeds by induction; it may help to think about agent i as $i \equiv i_0$.

In period $s = 1$, select an agent with the following property:

$$i_1 \in \arg\max_{j \in N} d(i, j; g(0))$$

where $d(i, j; g(0))$ refers to the geodesic distance (number of links along the shortest path) from i to j. Notice that by the definition of i_1, agent i's reach is independent from i_1's strategy, holding $g_{-i_1}(0)$ fixed.

It is now established that $g_{i_1} = \{i\}$ is a best response for i_1 to $g_{-i_1}(0)$. First, it is shown that this strategy yields a weakly higher payoff than any strategy with at least one link. For conciseness, let $g = \{g_{i_1} = \{i\}, g_{-i_1}(0)\}$. Consider the network $g' = (g'_{i_1}, g_{-i_1}(0))$ for any $g'_{i_1} \in G_{i_1}$. Let $N(i_1; g(0))$ be the set of agents who can reach i_1 in $g(0)$. Observe that:

$$\begin{aligned}
u_{i_1}(g) &= \sum_{j \in N(i_1; g(0))} \Phi\left(n, n - r_j(g(0)) + \Delta_{i_1 j}(g(0))\right) - C(1) \\
&= \sum_{j \in N(i_1; g(0))} \Phi\left(r_j(g') + [n - r_j(g')], [\Delta_{i_1 j}(g(0)) + r_j(g') - r_j(g(0))] + [n - r_j(g')]\right) - C(1) \\
&= \sum_{j \in N(i_1; g(0))} \Phi\left(r_j(g') + [n - r_j(g')], \Delta_{i_1 j}(g') + [n - r_j(g')]\right) - C(1) \\
&\geq \sum_{j \in N(i_1; g(0))} \Phi\left(r_j(g'), \Delta_{i_1 j}(g')\right) - C(1) \\
&\geq \sum_{j \in N(i_1; g(0))} \Phi\left(r_j(g'), \Delta_{i_1 j}(g')\right) - C(|g'_{i_1}|) = u_{i_1}(g')
\end{aligned}$$

The first equality owes to our previous argument that i's reach is independent from i_1's strategy, thus, $r_i(g) = n$. Thus, after that i_1 plays $\{i\}$, all agents in $N(i_1; g(0))$ can reach i and, therefore, have a reach of n. Any other agent's reach is not affected since they cannot reach i_1. The right side of the second and third equalities are equivalent to the first line (indeed, $\Delta_{i_1 j}(g') = \Delta_{i_1 j}(g(0)) + r_j(g') - r_j(g(0))$ for all $j \in N$). On the fourth line, the sum is a generic expression for the benefit that accrues to i_1 from any of their strategy in G_{i_1}, holding the others' strategies $g_{-i_1}(0)$ fixed. Assumption 1 part 3 permits to establish the first weak inequality on the fourth line. The second weak inequality owes to the fact that we are currently restricting our attention to strategies with links (so that $|g'_{i_1}| \geq 1$), and C is increasing.

Second, it is shown that $g_{i_1} = \{i\}$ yields a weakly higher payoff than the strategy of no links. This is immediate from the condition on the cost in Proposition 6, which states that having no links at all is a weakly dominated strategy.

This finishes to prove that forming one link to i is a best response for i_1 to $g_{-i_1}(0)$. Before moving on to the next period, it is useful to remind ourselves that agent i's reach is still equal to n in the network $g(1) = (g_{i_1} = \{i\}, g_{-i_1}(0))$ that is reached at the end of period 0.

In any subsequent period s such that $2 \leq s \leq n-1$, let $g(s-1)$ be the following network: there are $s-1$ distinct agents $i_1, \ldots, i_{s-1}$ such that, for any $q \in \{1, \ldots, s-1\}$, $g_{i_q}(s) = \{i_{q-1}\}$; also, $i \equiv i_0$ and $g_{i_1}(s-1) = \{i\}$. Choose i_s such that:

$$i_s \in \arg \max_{j \in N \setminus \{i_1, \ldots, i_{s-1}\}} d(i, j; g(s)).$$

Note that in $g(s-1)$, each agent in $i, i_1, \ldots, i_{s-1}$ can reach all n agents, and their reach is independent from i_s's strategy. To see why, note that for any $q \in \{i_1, \ldots, i_{s-1}\}$: (1) i_q reaches $i, i_1, \ldots, i_{q-1}$ via the link $i_q \to i_{q-1}$; (2) i_q reaches the agents in $N \setminus \{i, i_1, \ldots, i_{q-1}\}$ via agent i, whose reach is independent from i_s's strategy (since i_s is the furthest away from i and i already reaches the whole population). Consequently, similar reasoning as for period $s = 1$ applies to show that $g_{i_s} = \{i_{s-1}\}$ is a best response for agent i_s to $g_{-i_s}(s-1)$. The induction continues until i_{n-1} chooses a best response in the above manner. Hence, at the end of period $s = n-1$, agent $i \equiv i_0$ is the only agent with (eventually) more than one link. In the next period, let agent i move and play the following best response, $g_i = \{i_{n-1}\}$. The resulting network is then an exhaustive wheel. By Proposition 4, an exhaustive wheel is strict Nash for the cost considered and therefore absorbing under best-response dynamics.

It remains to prove that assuming $r_i(g(0)) = n$ for some $i \in N$ is without loss of generality. In period $s = 1$, select an agent with the following property:

$$i_1 \in \arg \max_{j \in N} r_j(g(0)).$$

By construction, the component to which i_1 belongs, referred to as A_1 below, is a maximal element of $\hookrightarrow$ given $g(0)$. Let i_1 play a best response to $g_{-i_1}(0)$, and let $g(1)$ be the resulting network. Note that we can consider that $|g_{i_1}(1)| \geq 1$, since the condition on the cost in Proposition 6 entails that having no link at all is a weakly dominated strategy.

Let i_1^* be an agent among those with the highest reach in $g(1)$. By the definition of $i_1, i_1^* \in A_1$. Let $R(i_1^*; g(1))$ be the set of agents whom i_1^* can reach in $g(1)$. If this set is N, we are done. Otherwise, we move on to the next period.

In period 2, select agent i_2 according to the following:

$$i_2 \in \arg \max_{j \notin R(i_1^*; g(1))} r_j(g(1)).$$

By construction, i_2 belongs to some component, called A_2 below, that is a maximal element of $\hookrightarrow$ given $g(1)$. Moreover, the constraint that $i_2 \notin R(i_1^*; g(1))$ entails that the strategy played by i_2 is irrelevant to i_1^*'s reach. As such, each agent in A_2 can reach at least $r_{i_1^*}(g(1)) + 1$ agents if i_2 forms a link with i_1^*.

Let i_2 play a best response to $g_{-i_2}(1)$, and let $g(2)$ be the resulting network. As argued in the previous period, it is without loss of generality that we can suppose that $|g_{i_2}(2)| \geq 1$. Also, the claim established below is that, for i_2^* an agent with the highest reach in $g(2)$, $r_{i_2^*}(g(2)) > r_{i_1^*}(g(1))$. To see why, assume by contradiction that $r_{i_2^*}(g(2)) \leq r_{i_1^*}(g(1))$. Hence, for each $j \in A_2$, $r_j(g(2)) \leq r_{i_1^*}(g(1))$. Note that the change in i_2's strategy between the two periods has only affected the reach of the agents in A_2. Because $r_j(g(1)) + [\Delta_{i_2 j}(g(2)) - \Delta_{i_2 j}(g(1))] = r_j(g(2)) \leq r_{i_1^*}(g(1))$ and $r_j(g(1)) = r_{i_2}(g(1))$, it follows that:

$$\delta_j^{i_2} := \Delta_{i_2 j}(g(2)) - \Delta_{i_2 j}(g(1)) \leq r_{i_1^*}(g(1)) - r_{i_2}(g(1)) := \delta^*. \tag{F5}$$

We will show that the strategy $g_{i_2} = \{i_1^*\}$ would have given a strictly higher payoff to i_2 than $g_{i_2}(2)$, which will imply the desired contradiction. To ease notations, set $g = (g_{i_2} = \{i_1^*\}, g_{-i_2}(1))$.

$$\begin{aligned}
u_{i_2}(g(2)) &= \sum_{j \in A_2} \Phi(r_j(g(2)), \Delta_{i_2 j}(g(2))) - C(|g_{i_2}(2)|) \\
&= \sum_{j \in A_2} \Phi(r_j(g(1)) + \delta_j^{i_2}, \Delta_{i_2 j}(g(1)) + \delta_j^{i_2}) - C(|g_{i_2}(2)|) \\
&< \sum_{j \in A_2} \Phi(r_j(g(1)) + \delta^* + 1, \Delta_{i_2 j}(g(1)) + \delta^* + 1) - C(1) \\
&\leq \sum_{j \in A_2} \Phi(r_j(g(1)) + [\Delta_{i_2 j}(g) - \Delta_{i_2 j}(g(1))], \Delta_{i_2 j}(g(1)) + [\Delta_{i_2 j}(g) - \Delta_{i_2 j}(g(1))]) - C(1) \\
&= u_{i_2}(g).
\end{aligned}$$

The strict inequality owes to (F5), Assumption 1 part 3, $|g_{i_2}(2)| \geq 1$ and the assumption that C is increasing. The argument that supports the weak inequality is as follows. In g, agent i_2 earns a benefit from the agents in A_2 only, and for any $j \in A_2$, this benefit can be written as:

$$\Phi(r_j(g), \Delta_{i_2 j}(g)) = \Phi(r_j(g(1)) + [\Delta_{i_2 j}(g) - \Delta_{i_2 j}(g(1))], \Delta_{i_2 j}(g(1)) + [\Delta_{i_2 j}(g) - \Delta_{i_2 j}(g(1))]).$$

Notice that, for each $j \in A_2$, $\Delta_{i_2 j}(g) = r_{i_1^*}(g(1)) + 1$ because both i_2 and j belong to the same component B in $g(1)$ while i_1^* does not belong to B in $g(1)$; further, in g, $g_{i_2} = \{i_1^*\}$. Also, $\Delta_{i_2 j}(g(1)) \leq r_{i_2}(g(1))$ because by construction, $\Delta_{i_2 j}(g(1)) \leq r_j(g(1))$ and since i_2, j belong to the same component B in $g(1)$, $r_j(g(1)) = r_{i_2}(g(1))$. Thus, $\Delta_{i_2 j}(g) - \Delta_{i_2 j}(g(1)) \geq \delta^* + 1$. It follows from Assumption 1 part 3 that:

$$\Phi(r_j(g), \Delta_{i_2 j}(g)) \geq \Phi(r_j(g(1)) + \delta^* + 1, \Delta_{i_2 j}(g(1)) + \delta^* + 1) \quad \forall j \in A_2.$$

Since the set of networks is finite and every best response is selected with positive probability, the finite paths constructed above occur with positive probability after any history. Hence no set of networks other than the exhaustive wheels can constitute a recurrent class of the dynamic process. Exhaustive wheels are absorbing states because, given the condition on the cost in Proposition 6, they are strict Nash. Therefore, the dynamic process converges to an exhaustive wheel with probability 1.

Appendix G Proof of Proposition 7

Consider a network g in which a link was formed by some agent i towards some other agent j, $i, j \in A_0$, and the link does not exist in g_0. Agent i has a strictly profitable deviation in cutting the link with j. This is because A_0 remains a component after the link is cut. Hence, the marginal benefit of the link $i \to j$ is null, while it is costly to form and maintain. Thus, if some network g is strict Nash, all links in g between the agents in A_0 must exist in g_0.

Consider a network g in which there is a link $i \to j$, where $i \notin A_0$ and $j \in A_0$. Since $|A_0| \geq 2$, there is an agent $k \in A_0$ such that $k \neq j$. If $k \notin g_i$, the deviation for player i, $g_i' = (g_i \setminus \{j\}) \cup \{k\}$, is payoff-equivalent to agent i's current strategy in g. If $k \in g_i$, the deviation $g_i' = g_i \setminus \{k\}$ is strictly profitable (same benefit, lower cost of link formation). Therefore, g is not strict Nash.

The previous argument establishes that A_0 remains a component in any strict Nash network. Then there are two cases worth distinguishing:

- **Case 1:** A_0 is both a maximal and minimal element of $\hookrightarrow$.
- **Case 2:** A_0 is maximal, not minimal.

Case 1: Consider the links formed by the agents outside A_0: $L(g) = \{i \to j | i \notin A_0\}$. By the previous argument, if g is strict Nash, $j \notin A_0$ for all links $i \to j \in L(g)$. We can use the previous analysis in Proposition 3. If g is strict Nash, the sub-network $g \setminus A_0$ whose set of nodes is $N \setminus A_0$ and whose set of links is $L(g)$ must have one of the following architectures: (1) the empty network, (2) an exhaustive wheel, (3) a non-exhaustive wheel, (4) an exhaustive tree, either with intermediaries or root-leaves.

First, I show that if g is strict Nash, the sub-network $g \setminus A_0$ cannot be a tree. Suppose instead that the agents in $g \setminus A_0$ form a tree. Hence, some agent $i \in N \setminus A_0$ currently has a link with some leaf j in the tree. The benefit from that link is to provide some agents in the tree access to leaf j. Let i deviate by cutting the link with j and adding a link with any agent $k \in A_0$. It is immediate from $|A_0| \geq 2$ and the fact that A_0 and the tree are not comparable via $\hookrightarrow$ that the deviation is strictly profitable. Indeed, the deviation entails the same cost of link formation, and it yields a strictly higher network benefit to agent i: agent i provides the same set of agents (namely, those who can reach i in the tree) access to strictly more individuals (from 1 agent, the leaf, to $|A_0|$ agents). Thus, g is not strict Nash.

It remains to show that if the sub-network $g \setminus A_0$ admits a wheel component A_1, $|A_1| > |A_0|$ if g is strict Nash. Consider instead the case where $|A_1| \leq |A_0|$. Consider a deviation for any agent $i \in A_1$ that consists of cutting the link in the wheel and simultaneously adding a link towards any agent $k \in A_0$. The same logic as above applies. Because $|A_0| \geq |A_1|$, and A_0 and A_1 are not comparable via $\hookrightarrow$, the deviation by the wheel member allows the same set of agents (A_1) to access strictly more individuals. Therefore, the deviation strictly increases the wheel member's benefit and maintains their cost of links formation. The deviation is then strictly profitable, hence the network is not strict Nash.

Case 2: Fix any network g in which A_0 is maximal, not minimal. If g is strict Nash, then:

1. Any component A such that $A_0 \hookrightarrow A$ is a singleton (by Corollary 2 Statement 1),

2. And any component A such that $A_0 \not\rightarrow A$ is an isolated singleton (by Corollary 2 Statement 3 and the assumption that A_0 is maximal).

I show here that the configuration in 2 above cannot occur in a strict Nash network. This is because an agent who maintains a link with a leaf in g has a payoff-equivalent deviation: to cut the link with the leaf and simultaneously add a link to an isolated singleton, which exists by configuration 2 above. Thus, in Case 2, a strict Nash network is an exhaustive tree whose root component is A_0.

Appendix H Proof of Corollary 11

Consider the version of the current game where link formation is cooperative as explained in Section 4.3, and the individual payoffs are given by (4). It was established in the discussion preceding the Corollary that an SBE is connected acyclic or the empty network.

The first part of the proof shows that a star network is the unique connected acyclic SBE. It is worthwhile remembering that by assumption, $n \geq 5$. See the Model Section 2.1. Consider any connected acyclic network g other than a star. Then, there always exist i, j, that satisfy the following five conditions:

1. $i \neq j$,
2. i is a leaf in g (i.e., $|g_i| = 1$)
3. j is not a leaf in g (i.e., $|g_j| \geq 2$),
4. there is no link between i and j in g,
5. For $g_i = \{k\}$, there exists $l \in g_k$ such that $l \notin \{i, j\}$. Meaning, if k is the agent with whom i has a link in g, agent k has a link with an agent other than i and j.

A few comments are in order. About conditions 3 and 4: since g is not a star, g is connected, the diameter of g (i.e., the length of the longest path in g) is at least 3. Hence, it is always possible to find an agent j two-links away from a leaf i. About condition 5: this condition is not satisfied if, and only if, $g_k = \{i, j\}$. Suppose this is the case; thus, in g, there is the following sequence of links: $j - k - i$. Because $n \geq 5$, j is not a leaf and g is acyclic, there exists $m \in g_j$ such that $m \neq i, j, k$. Hence, the following sequence of links exists in g: $m - j - k - i$. Since $n \geq 5$ and g is acyclic, either m is not leaf or j has a link with some agent p s.t. $p \notin \{m, k, j, i\}$. If m is not a leaf, the pair (i, m) satisfies all conditions 1 to 5 (i.e., in the conditions, replace j by m). If m is a leaf, agent j has a link with some agent p s.t. $p \notin \{i, m, k, j\}$ (because $n \geq 5$). Here, the pair (m, k) satisfies all five conditions (in the statements, replace i by m, j by k).

Given two agents i, j that satisfy all conditions in 1 to 5, consider the following deviation: for $m \in g_j$ such that m is crucial for connecting i and j in g, let j deviate by cutting the link with m and adding simultaneously a link with i. The deviation keeps the number of j's links constant, and it grants the same benefit to j. Indeed, if the path between i and j was $j - a_1 - \ldots - a_p$ where $a_1 \equiv m$, $a_p \equiv i$, the following path exists in the network after the deviation: $j - i - a_{p-1} - \ldots - a_1$. Thus, the deviation is payoff-equivalent to j. It remains to verify whether i accepts the link proposal. For agent i, accepting the link increases the cost of link formation by c. In the original network g, let:

- K be the set of agents who can no longer reach j if j cuts the link with m. Let $\kappa = |K|$ be the cardinality of that set. Note that $\{a_1, \ldots, a_{p-1}, i\} \subseteq K$ where $a_1 \equiv m$.
- Z be the set of agents who can no longer reach j if j cuts all of his links with the agents in $g_j \setminus \{m\}$. Let $z = |Z|$ be the cardinality of that set. Note that $Z = N \setminus (K \cup \{j\})$.

In the resulting network, agent i is crucial for connecting each agent in $K \setminus \{i\}$ to the agents in $Z \cup \{j\}$. The incremental benefit to agent i from accepting j's link proposal merely amounts to the cardinality of $(K \setminus \{i\}) \times (Z \cup \{j\})$, which is:

$$(\kappa - 1) \times (z + 1).$$

Therefore, i rejects the link proposal only if $c \geq (\kappa - 1)(z + 1)$.

Going back to the original network g, consider the deviation for leaf i whereby i cuts his link. This deviation is not profitable only if $c \leq n - 1$.

Gathering, none of the two deviations studied are profitable to agent i only if:

$$n - 1 \geq c \geq (\kappa - 1) \times (z + 1).$$

These two inequalities cannot hold simultaneously. First, notice that $z + \kappa = n - 1$. Second, $z \geq 1$. This is because j is not a leaf. Thus, j has a link with an agent other than m. Third, $\kappa \geq 3$; the reason is as follows. Since $i \notin g_j$ and $m \neq i, j$, one concludes that $\kappa \geq 2$. If $\kappa = 2$, so that $K = \{m, i\}$, condition 5 for the construction of the pair (i, j) is violated, as $K = \{m, i\}$ implies that $g_m = \{i, j\}$ (in condition 5, replace k by m). One can thus conclude that an acyclic connected SBE network is a star network.

In this modified version of the game where the payoffs are given by (4), an SBE is either a star network or the empty network. In the empty network, the marginal benefit of forming a link is 1. This network is an SBE if, and only if, $c > 1$.

Consider a star network. To mirror Goyal and Vega-Redondo (2007)'s Theorem 1, I seek to determine values for the cost of a link, c, for which no agents in a star want to form a link. Two leaves cannot gain from a deviation that adds a link between them. Also, two leaves cannot gain from cutting their link with the center and adding a link between them. Thus, the remaining deviations have two leaves forming a link between them, and one of them cuts the link with the center. Thus, take a leaf in a star, i. Let i cut the link with the center and add a link with another leaf, j. As argued in the first part of the proof, the deviation is payoff-equivalent for agent i. Consider agent j; if he accepts, j has to pay c to establish the link with i. The incremental benefit to agent j from the link with i is $n - 2$, as j is now crucial to all other individuals for reaching i. It follows that the link $i - j$ strictly decreases j's payoff if, and only if, $c > n - 2$.

Funding Open access funding provided by Corvinus University of Budapest.

References

Bala, V., Goyal, S.: A noncooperative model of network formation. Econometrica: Journal of the Econometric Society **68**(5), 1181–1229 (2000)

Billand, P., Bravard, C., Sarangi, S., Wooders, M.: The impact of price dispersion on R&D network formation. *Economic Theory*: https://doi.org/10.1007/s00199-026-01710-7 (2026)

Choi, S., Galeotti, A., Goyal, S.: Trading in networks: theory and experiments. J. Eur. Econ. Assoc. **15**(4), 784–817 (2017)

Feri, F.: Stochastic stability in networks with decay. J. Econ. Theory. **135**(1), 442–457 (2007)

Galeotti, A., Goyal, S., Kamphorst, J.: Network formation with heterogeneous players. Games Econom. Behav. **54**(2), 353–372 (2006)

Goyal, S., Vega-Redondo, F.: Structural holes in social networks. J. Econ. Theory. **137**(1), 460–492 (2007)

Hojman, D.A., Szeidl, A.: Core and periphery in networks. J. Econ. Theory. **139**(1), 295–309 (2008)

Jackson, M.O., Wolinsky, A.: A strategic model of social and economic networks. J. Econ. Theory. **71**(1), 44–74 (1996)

Joshi, S., Mahmud, A.S., Sarangi, S.: Network formation with multigraphs and strategic complementarities. J. Econ. Theory. **188**, 105033 (2020)

Kleinberg, J., Suri, É. Tardos, S., Wexler, T.: Strategic network formation with structural holes. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 284–293 (2008)

Li, B., Hao, D., Gao, H., Zhao, D.: Diffusion auction design. Artif. Intell. **303**, 103631 (2022)

Manea, M.: Intermediation and resale in networks. J. Polit. Econ. **126**(3), 1250–1301 (2018)

Olaizola, N., Valenciano, F.: A unifying model of strategic network formation. Internat. J. Game Theory **47**, 1033–1063 (2018)

Padilla, A.J., Pagano, M.: Sharing default information as a borrower discipline device. Eur. Econ. Rev. **44**(10), 1951–1980 (2000)

Siedlarek, J.P.: Intermediation in networks. Econ. Theor. **79**(3), 1083–1105 (2025)