Organizing Modular Production*

Niko Matouschek
Northwestern University

Michael Powell
Northwestern University

Bryony Reich
Northwestern University

January 21, 2022

Abstract

We characterize the optimal communication network in a firm with a modular production function, which we model as a network of decisions with a non-overlapping community structure. Optimal communication is characterized by two hierarchies that determine whom each agent receives information from and sends information to. Receiver rank depends only on module cohesion while sender rank also depends on decision-specific values of adaptation. When the hierarchies are the reverse of each other, optimal communication is bottom up in aggregate, and when they are the same, it has a core-periphery structure, in which the core contains the most cohesive modules.

Keywords: organization, communication, hierarchies, network design

JEL classifications: D23, D85, L23

*We thank Heski Bar-Isaac, Hector Chade, Krishna Dasaratha, Wouter Dessein, Itay Fainmesser, Ben Golub, Matt Jackson, Willemien Kets, Luis Rayo, Jeroen Swinkels, Alireza Tahbaz-Salehi, Eduard Talamàs, Yiqing Xing and participants at various conferences and seminars for comments and suggestions. We are especially grateful to Andrea Prat for discussing this paper at the NBER Organizational Economics Fall 2021 meeting and to Hossein Alidaee and Alexander Boche for excellent research assistance. Matouschek: n-matouschek@kellogg.northwestern.edu; Powell: mike-powell@kellogg.northwestern.edu; Reich: bryony.reich@kellogg.northwestern.edu.

1 Introduction

Over the last 60 years, the economy has experienced a sharp shift towards modular production (Baldwin and Clark 2000). Nowadays so many products are made by assembling separately produced modules that the 21st century has been called the *Modular Age* (Garud, Kumaraswamy, and Langlois 2009). The rise of modular production has the potential to change the organization of firms, the structure of industries, and the location of production. In this paper we take a first step towards exploring the economic implications of modular production by examining its impact on the internal organization of firms.

Herbert Simon anticipated the rise of modular production in 1962, when he observed that complex social, technological, and biological systems—large firms, mechanical watches, the human body—tend to be made up of communities or modules, groups of elements with stronger within than across group interactions (Simon 1962). The advantage of this modular structure, he argued, is that it allows systems to adapt to changes in the environment by making adjustments in a limited number of modules while leaving the rest of the system unchanged. The prevalence of modular structures has since been confirmed by the literature on community detection, which has documented them in a wide variety of contexts from the internet to the global air transportation network and the brain (Guimera et al. 2005, Meunier et al. 2009, and Fortunato 2010).

Two years after Simon published his article, IBM announced the first modular computer, the System/360. Until then computers had been tightly integrated systems of their constituent parts. A change in the processor or any other critical component required the design of an entirely new computer. This made it difficult to adopt new technologies and adapt computers to the idiosyncratic demands of different customers. The System/360 was designed to change all this. Its modular structure was a deliberate choice by IBM's executives who tasked their engineers with developing a computer that was made up of a small number of easily assemblable and exchangeable modules. Henceforth, when a supplier developed a better disk drive, or a customer needed more storage, IBM was able to adapt quickly. Not only did this make the System/360 an enormous financial success, it also changed how computers have been built ever since (Baldwin and Clark 1997 and 2000).

The move towards modular production has not been confined to the computer industry. Over the last few decades, firms across a wide range of industries followed in IBM's footsteps and developed products with modular production functions. Smartphones, airplanes, and electric cars are all made by assembling a limited number of modules. Even homes are now routinely assembled

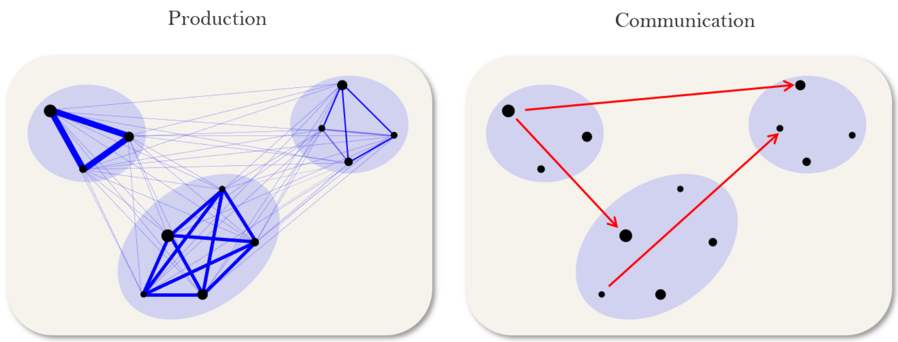

Figure 1: Left panel—The firm's production function takes the form of a network with a non-overlapping community structure. Right panel—Given the production network, the principal designs the optimal communication network by deciding whom each agent should tell his state to, taking as given that each directed link comes at an exogenous cost.

from pre-made modules rather than built on site from scratch. Nor is this move towards modular production confined to physical products. Modular programming—separating the different functions of computer programs into independent and interchangeable modules—was developed in the 1960s and is now pervasive. Of course, many products exhibited some degree of modularity even before the System/360. Builders installed pre-made doors and windows long before the rise of modular home building. What is different now, though, is that many products are modular by design. They are produced entirely by assembling a limited number of modules and, in line with Simon's observations, they are now more rule than exception.[1]

The widespread adoption of modular production has the potential to change the organization of production and, through this channel, the outcomes of economic activity. In the short run, firms adapt their internal organizations to accommodate modular production. Over time, they may also change their boundaries which, in turn, can alter the structure of their industries and the location of production. To manage the System/360, for instance, IBM established a centralized office, which ensured that different modules worked together, but also delegated control over individual modules to autonomous teams. This process of decentralization continued over many years with IBM and its competitors eventually outsourcing the development and production of modules to smaller, independent, and often foreign firms (Baldwin and Clark 1997 and 2000).

In this paper, we take a first step towards exploring the economic implications of modular production ("first step" in economics; there is an expansive existing literature in other fields which

[1] See Baldwin and Clark (1997, 2000). See also the Wikipedia entries for Modular Design, Modular Programming, and Modular Building and the references therein.

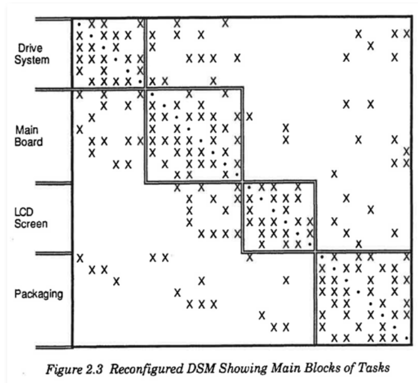

Figure 2: The "*Design Structure Matrix*" of a laptop computer in which each row and column corresponds to a task involved in producing the computer and an "x" entry indicates a strong need for coordination between the corresponding tasks (replication of Figure 2.3 in McCord and Eppinger (1993)).

we discuss below). We focus on the internal organization of a single firm with a modular production function, which we model as a network of decisions with a non-overlapping community structure as illustrated in Figure 1. Every node represents a decision, an agent who makes the decision, and a state that captures the local conditions. The size of a node represents the importance of adapting the decision to its local conditions, and the width of an edge represents the importance of coordinating the two decisions it connects. Decisions are partitioned into modules, groups of decisions that require more coordination with each other than with decisions in other modules. They are indicated by the shaded areas in the figure. The adjacency matrix of the production network, thus, takes the form of a block matrix. This structure approximates the interactions between decisions involved in making modular products, such as those for the laptop computer illustrated in Figure 2.

The only impediment to efficient production is that agents do not observe each other's local conditions. To improve efficiency, the principal can establish communication channels between agents. An IBM engineer who works on the disk drive does not directly observe the factors relevant to someone who works on the processor. But IBM can require the former to meet with the latter and learn about those factors. The problem is that such communication does not come for free. Even in the age of ever-evolving communication technologies, explaining the issues one faces to others takes time and energy, especially when they do not share the same expertise and experience. Given this trade-off between the efficiency of decision-making and the cost of communication, the principal decides whom each agent should tell his state to. In terms of Figure 1, the principal takes as given the production network in the left panel and designs the optimal communication network on the right by placing directed links between agents, taking into account that each link comes at an exogenous cost.[2]

The challenge in designing an optimal communication network is the abundance of possibilities and absence of any apparent way to order them. To see how the principal can overcome this challenge, it is useful to start by asking when she should add a single, directed link to an existing communication network. The cost of establishing such a link is the time and energy it takes the sender to explain his local conditions to the receiver, which we take to be exogenous. The benefit is that learning the sender's state allows the receiver to coordinate his decision more closely with the sender's, which, in turn, allows the sender to adapt his decision more closely to his state. Crucially, we find this benefit is independent of what the receiver, or any other agent, knows about any other state. Because of this independence, the problem of designing an optimal communication network can be separated into independent subproblems. It is sufficient for the principal to consider each agent in turn and ask whom this agent should tell about his state.

This separability result allows us to fully characterize optimal communication networks, which is our central result. Optimal communication takes the form of two hierarchies, one that determines whom each agent tells his state to and another that determines whose states he is told about. An agent with a higher sender rank tells his state to all the agents in other modules that a lower-ranked agent does, and possibly others. And an agent with a higher receiver rank is told about all the states in other modules that a lower-ranked agent is told about, and possibly more.

An agent's rank in either hierarchy depends critically on the *cohesion* of his module, which captures how distinct his module is from the rest of the production network. As such, it is increasing in the number of decisions that are in the module and the need for coordination between them, and decreasing in the *degree of coupling*, the need for coordination across modules. The more cohesive an agent's module is, the more important it is that the agent learns about the local conditions in other modules and that agents working on other modules learn about his.

Receiver rank is fully determined by module cohesion. Agents working on the same module have the same receiver rank; they either all learn a state or none of them do. The same is not true for sender rank, which can vary across agents working on the same module. The reason is that the benefit of sending information depends, not just on module cohesion, but also on the variability of

[2] Our focus on the trade-off between the efficiency of decision-making and the cost of communication is in line with the discussion of the design of optimal communication within firms in Arrow (1974, p.177): "*Since information is costly, it is clearly optimal, in general, to reduce the internal transmission...That is, it pays to have some loss in value for the choice of terminal act in order to economize on internal communication channels. The optimal choice of internal communication structures is a vastly difficult question.*"

the sender's local conditions and the need to adapt to them. If an agent's module is very cohesive, it is important he learns about other modules to enable them to adapt their decisions to their local conditions. But if his own local conditions are predictable, or the need to adapt to them is low, it may not be important for those working on other modules to learn about them. Sender and receiver ranks need not be perfectly correlated. Agents who hear a lot may speak little.

Different correlations between the optimal sender and the receiver hierarchies correspond to communication networks with different properties. If the optimal sender and receiver hierarchies are the reverse of each other, communication is bottom up in aggregate, involving more communication from agents in less cohesive modules to those in more cohesive ones than the reverse. If, instead, the optimal sender and receiver hierarchies are identical, the communication network has a core-periphery structure, in which agents in the most cohesive modules form the core and engage in intense communication with each other, while those in less cohesive modules make up the periphery and communicate only with the core and not with others in the periphery. Such core-periphery structures are pervasive among social and communication networks (see, for instance, Borgatti and Everett (2000) and Rombach et al. (2017) and the references therein).

After deriving the characterization result, we apply it to the notion that communication links should simply mirror technological interdependency, that "*we should expect to see a very close relationship...between a network graph of technical dependencies within a complex system and network graphs of organizational ties showing communication channels*" (Colfer and Baldwin 2016, p.713). This notion has a long history in management and related fields, where it is known as the *Mirroring Hypothesis* (see Thompson (1967) and, for a discussion of the literature, Colfer and Baldwin (2016)). In our setting, mirroring corresponds to a corner solution of the principal's problem that is only optimal if there are not too many modules and there is no single module that involves too many decisions or requires too much coordination.

At last, we return to Herbert Simon, who did not stop at observing that complex systems tend to be made up of modules. Rather, he noted that they often have, what he called, a *nearly decomposable structure*, in which the modules themselves are made up of sub-modules, sub-modules are made up of sub-sub-modules, and so on. In our extension, we allow for a symmetric version of a production function with such a nested, modular structure and show that optimal communication again follows a simple threshold rule.

2 Related Literature

To the best of our knowledge, there is no literature in economics on modular production. From a technical perspective, our paper belongs to the small but growing literature on centralized network design. In an early paper in this literature, Baccara and Bar-Isaac (2008) explore the optimal design of a network among members of a criminal organization in which more links facilitate cooperation but also leave the organization more vulnerable to attack by law enforcement. The trade-off between the efficiency of interactions among members of a network and its increased vulnerability to attacks by outsiders is also at the center of Goyal and Vigier (2014), who are motivated by the optimal design and defense of computer networks. Both papers differ from ours not only in terms of motivation but also modeling.

Closer to us is Calvó-Armengol and de Martí Beltran (2009). They consider an organization in which each agent's payoff depends on how well his decision is adapted to a common state, about which he is imperfectly informed, and coordinated with the other agents' decisions. A key feature of their model is that the production network is complete, with the need for coordination between any two decisions being the same, which precludes production from being modular. In this setting, they allow the principal to design the communication network, assuming she can add communication links at no cost up to an exogenously given cap. They show that, if the need for coordination is sufficiently small, and the degree of uncertainty sufficiently high, an optimal network maximizes a span index that they define.

Herskovic and Ramos (2020) also consider a setting in which each agent's payoff depends on how well his decision is adapted to a common state and coordinated with the other decisions. The production network is again complete, and thus production not modular. The key difference between their model and both Calvó-Armengol and de Martí Beltran (2009) and ours is that the communication network is not designed by a principal but formed by the agents' decentralized decisions of whom to communicate with. Their paper, therefore, belongs to the large literature on endogenous network formation that started with Jackson and Wolinsky (1996) and Bala and Goyal (2000), rather than the literature on centralized network design that ours belongs to. They show that even though agents' decisions are identical ex ante, the network they form is hierarchical, with agents in a given tier having their signals observed by those in the lower tiers.

In the endogenous network formation literature, the paper whose setting is closest to ours is Calvó-Armengol, de Martí, and Prat (2015). They, too, consider an organization whose agents face a trade-off between adaptation and coordination. Like us, though, they assume that each agent is adapting his decision to an independent state and, crucially, allow for decisions to differ in their needs for coordination. Even though they allow for differing coordination needs, however, they do not assume that production has a non-overlapping community structure, and thus do not explore modular production. Their main result characterizes how much effort each agent puts into both explaining his state to others and understanding theirs.

A feature shared by all the above papers on adaptation and coordination, and ours, is that the agents' payoff functions are quadratic, and their actions are continuous and exhibit strategic complementarities. As such, they all build on the literature on quadratic games on networks that started with Ballester, Calvó-Armengol, and Zenou (2006). In recent contributions to this literature, Bergemann, Heumann, and Morris (2017) and Golub and Morris (2017) characterize optimal decision-making for general information and network structures. We draw on their results to determine the agents' decision-making for given communication networks. A property of equilibrium decision-making in our setting is that it depends on the value of cyclical walks on the production network, which relates to the notion of cyclical centrality in Talamàs and Tamuz (2017). Our focus, though, is not on decision-making but on the prior stage in which the principal designs the communication network, taking as given that agents will make their decisions optimally.

Even though our modeling approach places us in the network literature, the primary literature our paper belongs to is the literature on organizational economics and, in particular, team theory. Starting in the 1950s, team theory explores the optimal design of organizations when agents share the same goal, but cognitive constraints make communication costly (for an early treatment see Marschak and Radner (1972) and for recent surveys see Garicano and Prat (2013) and Garicano and Van Zandt (2013)). A related paper in this literature is Dessein and Santos (2006), who were the first to explore how the trade-off between adaptation and coordination shapes the internal organization of firms. In their setting, the production network is complete, with the need for coordination between any two decisions being the same, and the principal does not design a communication network. Instead, they allow for each agent to make multiple decisions and assume the same quality of communication between any pair of agents. They show that more uncertainty about the environment increases the optimal number of decisions per agent, while the effect of an improvement in the quality of communication on specialization is non-monotonic.

We also relate to Dessein, Galeotti, and Santos (2016), who build on Dessein and Santos (2006) by endogenizing communication while taking the allocation of decisions as given. In their setting, the production network is once again complete, which precludes production from being modular. They show that if the total amount of time that agents have to learn about others is limited, the principal finds it optimal to have them spend all their time learning about a small number of core agents, while staying largely ignorant about the others.

Even though, to our knowledge, there is no literature on modularity in economics, there is a large literature on this topic in management and related fields, as well as in computer science. As noted earlier, Simon (1962) observed that complex systems are often made up of modules and argued that this modular design facilitates adaptation. A similar point was made by Alexander (1964), who argued that a modular system design accelerates adaptation by allowing the system to adapt module by module. In computer science, Parnas (1972) argued that a modular software design allows for faster programming by enabling different teams to work on different program modules in parallel and explored criteria to best decompose a program into modules.

Our paper connects to a related literature that takes the modular design of products as given and explores its implications for the organization of production. A common argument in this literature is the Mirroring Hypothesis we mentioned in the introduction, which posits that the organization of a firm, and specifically its internal communication structure, ought to mirror the modular nature of its production function. A firm that makes a modular product, in other words, should see intense communication within modules but not across (see, in particular, Thompson (1967), Henderson and Clark (1990), Sanchez and Mahoney (1996) and, for a survey, see Baldwin and Colfer (2016)). Langlois and Robertson (1992) observed that modular production might not only affect the internal organization of firms but also their boundaries and, through this channel, the structure of industries. Baldwin and Clark (2000) document these dynamics in the context of IBM and the computer industry, and provide an exhaustive discussion of modular production and its organization.

A related literature reverses the causality of the Mirroring Hypothesis and argues that the design of products reflects the organization of the firms that developed them. In this view, a modular organization has a tendency to develop modular products. In computer science, this view is known as *Conway's Law*, named after Melvin Conway who observed that "*To the extent that an organization is not completely flexible in its communication structure, that organization will stamp out an image of itself in every design it produces*" Conway (1968, p.30).

3 Model

A firm consists of one principal and N agents. All parties are risk neutral and care only about the firm's profits. There are no incentive conflicts.

Production. Each agent $i \in \mathcal{N}$ makes a decision $d_i \in [-D, D]$ that is associated with a state $\theta_i \in [-D, D]$, where $\mathcal{N} = \{1, ..., N\}$ is the set of agents, and D is a large but finite scalar. Output depends on how well each decision is adapted to its associated state and coordinated with the other decisions. Specifically, we follow Ballester, Calvó-Armengol, and Zenou (2006) and assume that output is given by

$$r\left(d_1,...,d_n\right)=\sum_{i=1}^{N}\left[-d_i^2+2a_id_i\theta_i+\sum_{j=1}^{N}p_{ij}d_id_j\right], \tag{1}$$

where $a_i>0$ captures the importance of adapting decision d_i to its state θ_i, and the degree of strategic complementarity $p_{ij}\geq 0$ captures the need for coordination between decisions d_i and d_j.[3] The need for coordination is symmetric, that is, $p_{ij}=p_{ji}$, and p_{ii} is equal to zero. The interactions between decisions can, therefore, be represented by an undirected network, which we summarize in an $N\times N$ matrix $\boldsymbol{P}$ with entries p_{ij}. We assume that $\sum_{j=1}^{N}p_{ij}<1$ for all $i\in\mathcal{N}$, which ensures that equilibrium decisions exist. Finally, we normalize the price of the product to one so that output (1) also represents revenue.

Modules. Each decision, and its associated state and agent, belongs to a "*module*" $\mathcal{M}_m$ for $m\in\{1,...,M\}$, which is a set of $n_m\geq 1$ such decisions. The function $m\left(i\right)$ gives the module $\mathcal{M}_{m(i)}$ that decision d_i belongs to. For expositional convenience we adopt the convention that the first decision d_1, and its associated state and agent, belong to module $\mathcal{M}_1$, and assume that there are at least three modules, that is, $M\geq 3$.

The need for coordination is stronger between two decisions within the same module than between two decisions in different modules. Specifically, the need for coordination between any two decisions d_i and $d_j\neq d_i$ is given by $p_{ij}=t\geq 0$ if they belong to different modules and, abusing notation slightly, it is given by $p_{ij}=p_m\geq t$ if they belong to the same module $\mathcal{M}_m$. The parameter t, therefore, captures the "*degree of coupling*" between modules, while the parameter p_m captures the need for coordination within module $\mathcal{M}_m$.

Information. Each agent $i\in\mathcal{N}$ privately observes the realization of his state θ_i, which is independently drawn from a distribution with zero mean and variance σ_i^2. All other information is public.

Before the states are drawn, the principal can place directed communication links between any

[3] This formulation permits as a special case the widely-used payoff function in which payoffs are the weighted average of the quadratic difference between each decision and its state and between each pair of decisions (see, for instance, Alonso, Dessein, and Matouschek (2008) and Calvó-Armengol and de Martí Beltran (2009)). Specifically, if $a_i=1-\sum_{j=1}^{N}p_{ij}$ for all $i\in\mathcal{N}$, we can re-write output (1) as

$$\sum_{i=1}^{N}\left[-\left(1-\sum_{j=1}^{N}p_{ij}\right)\left(d_i-\theta_i\right)^2-\frac{1}{2}\sum_{j=1}^{N}p_{i,j}\left(d_i-d_j\right)^2\right]+\sum_{i=1}^{N}\left(1-\sum_{j=1}^{N}p_{ij}\right)\theta_i^2,$$

where the last term is a constant.

two agents. Each such link comes at a cost, which captures the resources involved in communication. We assume that this cost is $\gamma > 0$ if the two agents belong to different modules, and it is zero if they belong to the same one. This assumption streamlines the exposition and is consistent with the notion that communication is less costly between agents who share similar expertise and experiences.

If the principal places a communication link from agent i to agent j, agent i tells j the realization of his state θ_i. Communication, therefore, takes the form of a directed network, which we summarize in an $N \times N$ matrix $\boldsymbol{C}$. Entry c_{ij} is equal to one if agent i tells agent $j \neq i$ his state and it is zero if he does not. Moreover, since each agent i observes his own state, c_{ii} is always equal to one. Row $\boldsymbol{C}_i$ then summarizes the agents who learn θ_i and column $\boldsymbol{C}_{(j)}$ summarizes the states agent j learns about.

Organization. The principal's problem is to design the optimal communication network that maximizes expected revenue net of communication costs, that is, to solve

$$\max_{\boldsymbol{C}} \mathrm{E}[r\left(d_1, ..., d_N\right) | \boldsymbol{C}] - \gamma \sum_{i=1}^{N} \sum_{j=1}^{N} m_{ij} c_{ij} \text{ subject to } c_{ii} = 1 \text{ for all } i \in \mathcal{N}, \tag{2}$$

where m_{ij} is a dummy variable that is equal to one if and only if agents i and $j \neq i$ belong to different modules.

Timing. After the principal designs the communication network, agents learn their states and tell them to other agents as specified by the network. Next, the agents simultaneously make their decisions, payoffs are realized, and the game ends. The solution concept we use is Perfect Bayesian Equilibrium.

We discuss the key assumptions, such as the assumption that agents do not re-transmit information they receive and that their decisions are not distorted by incentive conflicts, in Section 6, after we solve the model in the next section and apply it to the Mirroring Hypothesis in Section 5.

4 Solving the Model

To solve the model, we start by determining equilibrium decisions for any given communication network. We then show that given these equilibrium decision rules, we can simplify the principal's problem of designing an optimal communication network by separating it into independent subproblems. Finally, we characterize the solution to the principal's problem by solving these subproblems.

4.1 Decision-Making

After agents have observed their states and communicated with each other, they make the decisions that solve

$$\max_{d_i} \mathrm{E}\left[r\left(d_1, ..., d_N\right) \left| \boldsymbol{C}_{(i)}\right.\right] \text{ for all } i \in \mathcal{N},$$

where $r\left(d_1, ..., d_N\right)$ is revenue (1) and where $\boldsymbol{C}_{(i)}$ is the ith column of the communication matrix $\boldsymbol{C}$ that summarizes the states agent i is informed about. The best-response functions that follow from these optimization problems are given by

$$d_i = a_i \theta_i + \sum_{j=1}^{N} p_{ij} \mathrm{E}\left[d_j \left| \boldsymbol{C}_{(i)}\right.\right]. \tag{3}$$

Each agent's best response is the weighted sum of his state and the decisions he expects the other agents to make, where the weight on his own state is a_i and the weight on the decision he expects agent j to make is p_{ij}. To solve the system of best responses, note that $\left(\mathrm{diag}\boldsymbol{C}_j\right) \boldsymbol{P} \left(\mathrm{diag}\boldsymbol{C}_j\right)$ is the subgraph of the production network that consists of the nodes whose agents know state θ_j, as well as all the links between them. We can then state the following lemma.

LEMMA 1. *Equilibrium decisions are unique and given by*

$$d_i^* = \sum_{j=1}^{N} a_j \omega_{ij} \left(\boldsymbol{C}_j\right) \theta_j \text{ for all } i \in \mathcal{N}, \tag{4}$$

where $\omega_{ij}\left(\boldsymbol{C}_j\right)$ *denotes the* ij*th entry of* $\left(\boldsymbol{I} - \left(\mathrm{diag}\boldsymbol{C}_j\right) \boldsymbol{P} \left(\mathrm{diag}\boldsymbol{C}_j\right)\right)^{-1}$.

The lemma shows that agent i's equilibrium decision d_i^* is the weighted sum of all states, where the weight on state θ_j is given by a_j, the importance of adapting decision d_j to θ_j, times $\omega_{ij}\left(\boldsymbol{C}_j\right)$, the ijth entry of $\left(\boldsymbol{I} - \left(\mathrm{diag}\boldsymbol{C}_j\right) \boldsymbol{P} \left(\mathrm{diag}\boldsymbol{C}_j\right)\right)^{-1}$. This latter object has a natural interpretation in terms of walks on the production network. Before providing it, though, we pause briefly to review the notion of walks and their values.

A "*walk*" between d_i and d_j on the production network is a sequence of links that lead from d_i to d_j. Each link between two decisions in this sequence is associated with a discount factor, which is given by the need for coordination between them. The "*value of a walk*" is the product of these discount factors. As an example, consider the production network $\boldsymbol{P}^e$ in Figure 3, where the superscript stands for "example." In this case, d_1 to d_2 to d_3 constitutes a walk from d_1 to d_3 whose value is given by t^2. Standard arguments imply that the ijth entry of $\left(\boldsymbol{I} - \boldsymbol{P}^e\right)^{-1}$ is the sum of the values of all walks from d_i to d_j on the production network $\boldsymbol{P}^e$.

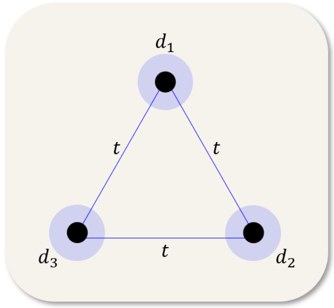

Figure 3: Production network $\boldsymbol{P}^e$, which consists of three modules that contain one decision each. The parameter t is the degree of coupling, the need for coordination across modules.

In light of this discussion, the ijth entry of $(\boldsymbol{I} - (\text{diag}\boldsymbol{C}_j)\,\boldsymbol{P}\,(\text{diag}\boldsymbol{C}_j))^{-1}$ in agent i's equilibrium decision rule (4) represents the value of all walks from d_i to d_j on the subgraph of the production network that consists only of decisions made by agents who know state θ_j. If agent i does not know θ_s, for instance, d_i is not part of this subgraph, and so $\omega_{is}(\boldsymbol{C}_s) = 0$. Agent i puts no weight on θ_s, as one would expect. If, instead, θ_s is public, the subgraph encompasses the entire production network, and the weight agent i puts on θ_s is the value of all walks from d_i to d_s on the production network $\boldsymbol{P}$. Note that this is the case no matter what the agents know about the other states. This result reflects a general implication of the lemma that will be important for what follows: for a given production network, the weight agent i puts on state θ_s depends only on who knows θ_s and not on what agent i, or any other agent, knows about any other state.

The part of the equilibrium decision rules that will turn out to be key for the optimal design of communication networks is the weight each decision puts on its own state. To get an intuition for this weight, recall the production network $\boldsymbol{P}^e$ in Figure 3, and suppose agent 1 does not tell his state to the other two agents. Agent 1 is then forced to adapt to his state autonomously, without the benefit of having the others coordinate their decisions with his. This limits the weight he puts on his own state to

$$a_1\omega_{11}^e\left((1,0,0)\right) = a_1,$$

where the superscript again stands for "example." The parameter a_1, therefore, captures the "*degree of autonomous adaptation*."

Suppose now that agent 1 tells his state to agent 2 but still not to agent 3. Since agent 2 cares about coordinating his decision with agent 1's, he will put some weight on θ_1. And since agent 1 also cares about coordinating his decision with agent 2's, this induces him to put more weight on his own state. Specifically, if agent 1 tells his state to agent 2, the weight agent 1 puts on θ_1

increases to

$$a_1 \omega_{11}^e \left((1,1,0) \right) = a_1 \left(1 + \frac{t^2}{1-t^2} \right) > a_1.$$

Communication, in other words, enables coordination, which, in turn, facilitates adaptation. The extent to which it does so is captured by "*coordination multiplier*" $\omega_{11} \left(\boldsymbol{C}_1 \right)$.

A key property of the coordination multiplier is that it is supermodular. Suppose agent 1 tells his state to agents 2 and 3. Since agents 2 and 3 care about coordinating with each other, and not just with agent 1, they will put more weight on θ_1 than they would if agent 1 told his state to only one of them. This increase in the weights agents 2 and 3 put on θ_1, in turn, induces agent 1 to increase the weight he puts on his own state to

$$a_1 \omega_{11}^e \left((1,1,1) \right) = a_1 \left(1 + \frac{t^2}{1-t^2} + \frac{t^2}{1-t^2} + \frac{2t^3}{\left(1-t^2\right)\left(1-2t\right)} \right),$$

where the last term in brackets captures the supermodularity.

These properties of the equilibrium decision rule hold in general, and we summarize them in the following corollary.

COROLLARY 1. *The weight $a_i \omega_{ii} \left(\boldsymbol{C}_i \right)$ that agent i's decision d_i^* puts on his state θ_i satisfies $\omega_{ii} \left(\boldsymbol{I}_i \right) a_i = a_i$, where $\boldsymbol{I}_i$ is* ith *row of an $N \times N$ identity matrix, and is increasing and supermodular in $\boldsymbol{C}_i$.*

Having characterized equilibrium decision-making by the agents, we next turn to the principal's problem.

4.2 Simplifying the Principal's Problem

The principal's problem is to design the communication network that maximizes expected profits, taking into account that agents make decisions according to (4). It is useful to start by rewriting revenue (1) as

$$r \left(d_1, ..., d_N \right) = \sum_{i=1}^{N} a_i d_i \theta_i - \sum_{i=1}^{N} d_i \left(d_i - a_i \theta_i - \sum_{j=1}^{N} p_{ij} d_j \right).$$

Substituting in the equilibrium decision rules (4), this simplifies to

$$r \left(d_1^*, ..., d_N^* \right) = \sum_{i=1}^{N} a_i d_i^* \theta_i + \sum_{i=1}^{N} \sum_{j=1}^{N} p_{ij} d_i^* \left(d_j^* - \mathrm{E} \left[d_j^* \left| \boldsymbol{C}_{(i)} \right. \right] \right). \tag{5}$$

In the proof of the next lemma we show that the second term on the right-hand side is zero in expectation, which implies the following result.

LEMMA 2. *Under equilibrium decision-making, expected revenue is given by*

$$R(\boldsymbol{C}) \equiv \mathrm{E}\left[r\left(d_1^*, ..., d_N^*\right)\right] = \sum_{i=1}^{N} a_i \mathrm{Cov}\left(d_i^*, \theta_i\right), \tag{6}$$

where $\mathrm{Cov}\left(d_i^*, \theta_i\right) = a_i \sigma_i^2 \omega_{ii}\left(\boldsymbol{C}_i\right)$.

The lemma shows that expected revenue boils down to how well each decision is adapted to its associated state. For expositional convenience, we interpret $a_i \mathrm{Cov}\left(d_i^*, \theta_i\right)$ as the expected revenue generated by agent $i \in \mathcal{N}$ and denote it by

$$R_i\left(\boldsymbol{C}_i\right) \equiv a_i \mathrm{Cov}\left(d_i^*, \theta_i\right) = a_i^2 \sigma_i^2 \omega_{ii}\left(\boldsymbol{C}_i\right).$$

The key property of agent i's expected revenue is that it depends on $\boldsymbol{C}_i$ but not on the rest of communication network $\boldsymbol{C}$. An additional agent learning θ_i increases agent i's coordination multiplier $\omega_{ii}\left(\boldsymbol{C}_i\right)$, and thus the weight $a_i \omega_{ii}\left(\boldsymbol{C}_i\right)$ he puts on his state, as well as the expected revenue $a_i^2 \sigma_i^2 \omega_{ii}\left(\boldsymbol{C}_i\right)$ he generates. In contrast, agent i, or any other agent, learning any other state does not affect $\omega_{ii}\left(\boldsymbol{C}_i\right)$, and thus leaves the weight agent i puts on his state, and the revenue he is expected to generate, unchanged.

This property of expected revenue is key because it implies that the principal's problem is separable. Instead of solving the overall problem (2) head on, the principal can consider each agent in isolation and ask whom this agent should tell about his state. The answer to whom agent $i \in \mathcal{N}$ should tell about θ_i is independent of whom any other agent should tell about his own state. We, therefore, have our first main result.

PROPOSITION 1. *An optimal communication network solves the principal's problem* (2) *if and only if it solves the* N *independent subproblems*

$$\max_{\boldsymbol{C}_i} R_i\left(\boldsymbol{C}_i\right) - \gamma \sum_{j=1}^{N} m_{ij} c_{ij} \text{ subject to } c_{ii} = 1 \text{ for all } i \in \mathcal{N}, \tag{7}$$

where m_{ij} *is a dummy variable that is equal to one if and only if agents* i *and* $j \neq i$ *belong to different modules.*

This separability result greatly facilitates the principal's quest for optimal communication networks. We can further simplify the problem by recalling that agent i's coordination multiplier $\omega_{ii}\left(\boldsymbol{C}_i\right)$ is supermodular. This property implies that, whenever it is optimal for agent i to tell agent j about his state, it must also be optimal for him to tell the other agents in agent j's module $\mathcal{M}_{m(j)}$. The principal's problem, therefore, reduces to which modules each agent should tell about his state.

Finally, supermodularity of $\omega_{ii}(\cdot)$, together with the linearity of the communication costs, imply that these subproblems are also supermodular. For any given parameter values, the principal's problem can, therefore, be solved using standard algorithms that maximize supermodular functions in polynomial time (see, for instance, chapter 10.2 in Murota (2003)). Our goal, though, is to solve the problem analytically, and we do so in the next section.

4.3 Optimal Communication Networks

The separability result in Proposition 1 allows us to solve the principal's problem by considering each agent in isolation and asking whom he should tell about his state. The first step in answering this question is to express the agent's expected revenue in terms of the model's primitives. To economize on notation, and without loss, we focus on agent 1.

Suppose agent 1 tells his state to all agents in an arbitrary set of modules that includes his own module $\mathcal{M}_1$. Since the naming of modules is immaterial, there is no loss in denoting the modules in this set by $\mathcal{M}_1$, ..., $\mathcal{M}_\ell$ for $\ell \in \{1, ..., M\}$. We can then define $\boldsymbol{C}_1(\ell)$ as the row of the communication matrix that specifies the agents who belong to these modules and, thus, know θ_1. The next lemma uses this notation to express agent 1's expected revenue in terms of model primitives.

LEMMA 3. *Suppose agent* 1 *tells his state to all agents in modules* $\mathcal{M}_1$, ..., $\mathcal{M}_\ell$, *for* $\ell \in \{1, ..., M\}$ *and to none of the agents in other modules. Agent* 1*'s expected revenue is then given by*

$$R_1(\boldsymbol{C}_1(\ell)) = a_1^2 \sigma_1^2 \left(\frac{1-(n_1-2)\,p_1}{(1+p_1)(1-(n_1-1)\,p_1)} + \frac{t^2 x_1^2 \left(\sum_{m=1}^{\ell} n_m x_m - n_1 x_1 \right)}{(1-t n_1 x_1)\left(1 - t \sum_{m=1}^{\ell} n_m x_m\right)} \right), \tag{8}$$

where

$$x_m \equiv \frac{1}{1-(n_m-1)\,p_m + n_m t} \text{ for } m = 1, ..., M.$$

The object x_m in the lemma is the "*cohesion*" of module $\mathcal{M}_m$, which captures how distinct the module is from the rest of the production network.[4] As such, it is increasing in its size and the need for coordination among its members and decreasing in the degree of coupling t.

The lemma shows that agent 1's expected revenue is the product of $a_1^2\sigma_1^2$ and the coordination multiplier $\omega_{11}(\boldsymbol{C}_1(\ell))$. The term $a_1^2\sigma_1^2$ is the revenue agent 1 is expected to generate if he does not tell his state to any other agent and adapts to his state autonomously. As such, we refer to $a_1^2\sigma_1^2$ as the "*value of autonomous adaptation*" of decision d_1. The coordination multiplier, in turn, is

[4] There are different notions and formal definitions of cohesion in the sociology and economics literatures. Our definition is close to that in Morris (2000). Applied to our setting, his definition of the cohesion of module $\mathcal{M}_m$ is $(n_m-1)\,p_m / [(n_m-1)\,p_m + (N-n_m)\,t]$.

the sum of two terms. The first is the value of all walks from node 1 back to itself that only go through nodes in $\mathcal{M}_1$. Because these walks only go through module $\mathcal{M}_1$, their value depends only on its characteristics p_1 and n_1. The second term, then, is the value of the additional walks that also go through at least one node in one of the other modules $\mathcal{M}_2$,..., $\mathcal{M}_\ell$. Notice that the value of these additional walks depends on the characteristics of $\mathcal{M}_1$ only through its cohesion x_1 and the scaled analog of its cohesion $n_1 x_1$, and that it depends on the characteristics of the other informed modules only through the sum of their $n_m x_m$ terms. Agent 1's expected revenue, for instance, is the same whether he tells his state to agents in one module with $n_2 x_2 = 10$ or to agents in ten modules with $n_2 x_2 = ... = n_{11} x_{11} = 1$. This property facilitates the characterization of optimal communication networks, to which we turn next.

PROPOSITION 2. *Optimal communication is characterized by N thresholds $\lambda_i \geq 0$, one for each agent $i \in \mathcal{N}$. Agent i tells agent j about his state if and only if they belong to the same module or the cohesion of agent j's module is above agent i's threshold, that is, $x_{m(j)} \geq \lambda_i$. The threshold λ_i is increasing in marginal communication costs γ and decreasing in the value of autonomous adaptation $a_i^2 \sigma_i^2$, the need to coordinate the decisions within agent i's module $p_{m(i)}$, and the size of his module $n_{m(i)}$.*

In an optimal communication network, an agent always tells his state to the other agents in his own module. This result follows immediately from the assumption that the cost of doing so is zero and the fact that the benefit is strictly positive. The key insight in the proposition is that, apart from the agents in his own module, an agent will tell his state to those in the most cohesive modules. Doing so increases the covariance between his decision and state by enough to warrant the additional communication costs. Telling his state to agents in less cohesive modules still allows him to adapt his decision more aggressively to his state, but not by enough to justify the additional costs.

The proposition is illustrated in Figure 4, in which we again focus on agent 1. There are five modules, and modules $\mathcal{M}_2, ..., \mathcal{M}_5$ are labeled in decreasing order of their cohesion, so that $x_2 \geq x_3 \geq x_4 \geq x_5$. The blue curve is the piecewise linear extension of expected revenue $R_1 (\boldsymbol{C}_1 (\ell))$, which we denote by $\overline{R}_1 (\boldsymbol{C}_1 (\ell))$, and the red line is a continuous representation of communication costs $\sum_{m=2}^{\ell} n_m \gamma$. The changing curvature of expected revenue $\overline{R}_1 (\boldsymbol{C}_1 (\ell))$ reflects the countervailing economic forces at work. The supermodularity at the heart of the model pushes towards convexity while the modular structure of the production function pushes towards concavity. A reduction in γ flattens the cost curve, which favors telling agents in more modules about θ_1. For the other comparative statics in the proposition, note that the slope of each line segment in the

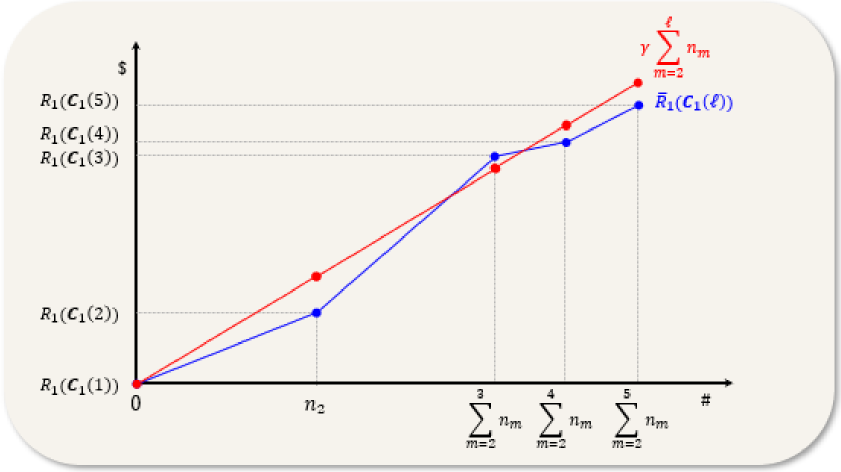

Figure 4: Determining the optimal communication network for agent 1 (drawn for parameter values $t = 0.01$, $n_1 = n_2 = n_3 = 5$, $n_4 = n_5 = 2$, $p_1 = p_2 = p_3 = 0.2$, $p_4 = p_5 = 0.1$, $a_1\sigma_1 = 1$).

benefits curve is the additional expected revenue generated by telling agents in the corresponding module about θ_1 divided by the number of agents in the module. From Lemma 3, this per node marginal benefit is given by

$$\frac{1}{n_{\ell+1}} \left(R_1 \left(\boldsymbol{C}_1 \left(\ell + 1 \right) \right) - R_1 \left(\boldsymbol{C}_1 \left(\ell \right) \right) \right) = a_1^2 \sigma_1^2 \frac{t^2 x_1^2 x_{\ell+1}}{\left(1 - t \sum_{m=1}^{\ell} n_m x_m \right) \left(1 - t \sum_{m=1}^{\ell+1} n_m x_m \right)},$$

which is increasing in the value of autonomous adaptation $a_1^2\sigma_1^2$, the size of agent 1's module n_1, and the need for coordination among its members p_1. An increase in any of these parameters, therefore, steepens the benefits curve, which favors telling agents in more modules about θ_1.

The proposition implies that optimal communication gives rise to sender and receiver hierarchies. To see this implication clearly, focus again on agent 1 and consider the agents in modules $\mathcal{M}_2$ and $\mathcal{M}_3$. The proposition shows that if $\mathcal{M}_2$ is more cohesive than $\mathcal{M}_3$, agents in module $\mathcal{M}_3$ will only ever be told about θ_1 if those in module $\mathcal{M}_2$ also are. Moreover, since x_2 and x_3 do not depend on the characteristics of the sender's module $\mathcal{M}_1$, agents in module $\mathcal{M}_3$ will only ever be told about *any* state in another module that those in module $\mathcal{M}_2$ also are. Optimal communication, therefore, gives rise to a *receiver hierarchy*, in which a higher-ranked agent is told about all the states in other modules that a lower-ranked agent is told about, and possibly more. Agent i's receiver rank is fully determined by the cohesion $x_{m(i)}$ of the module he belongs to. If $x_{m(i)} \geq x_{m(j)}$, agent i outranks agent j.

COROLLARY 2. *Optimal communication gives rise to a receiver hierarchy among agents. For any agents $i, j, k \in \mathcal{N}$ who belong to different modules, if $x_{m(i)} \geq x_{m(j)}$, then agent j is told about agent k's state only if agent i also is.*

Notice that this result is about communication and not information per se. A higher-ranked agent is told about all the states in other modules that a lower-ranked agent is told about. But a lower-ranked agent may still have some information that a higher-ranked agent does not have. In particular, a lower-ranked agent observes his own state, and is always told about the other states in his own module, and it may well be optimal for a higher-ranked agent to remain ignorant about those states.

The fact that cohesion x_m depends only on the characteristics of module $\mathcal{M}_m$, and not on those of any other modules, also implies a *sender hierarchy* in which a higher-ranked agent tells his state to all the agents in other modules that a lower-ranked agent does, and possibly others. Agent i's sender rank, though, does not depend on $x_{m(i)}$ but on λ_i. Agent i outranks agent j in the sender hierarchy if he has a lower threshold, that is, $\lambda_i \leq \lambda_j$.

COROLLARY 3. *Optimal communication gives rise to a sender hierarchy among agents. For any agents $i, j, k \in \mathcal{N}$ who belong to different modules, if $\lambda_i \leq \lambda_j$, then agent j tells agent k about his state only if agent i also does.*

Other things equal, an agent's sender rank is increasing in the cohesion of his module, just as his receiver rank is. Yet, his sender and receiver ranks need not coincide. Agent i may outrank agent j in one hierarchy but be outranked by him in the other. The reason is that while agent i's rank in the receiver hierarchy depends only on the cohesion of the modules, his rank in the sender hierarchy also depends on the specific characteristics of the decisions, as captured by the values of autonomous adaptation $a_i^2\sigma_i^2$. Suppose agent i's module is very cohesive, so that $x_{m(i)}$ is larger than the x_m of any other module. Agent i, and the other agents in his module, then reside on top of the receiver hierarchy, being told the states of any modules that agents in other modules are told about. They reside on top of the receiver hierarchy because their ignorance about other modules would hold those modules back from adapting their decisions more than the ignorance of agents in any other module would. At the same time, if $a_i^2\sigma_i^2$ is sufficiently small, λ_i is also larger than the λ_j of any other agent $j \in \mathcal{N} \backslash i$, placing agent i at the bottom of the sender hierarchy. Even though his module is very cohesive, his ability to adapt his decisions to his state is just not very important. The agents who hear the most, therefore, might also speak the least.

Different correlations between the optimal sender and the receiver hierarchies give rise to communication networks with different properties. To build on the illustration in the previous paragraph, suppose first that the optimal hierarchies are the reverse of each other, that is, for all $i, j \in \mathcal{N}$, agent i outranks agent j in the receiver hierarchy if and only if the reverse holds in the sender hierarchy. Loosely speaking, agents in more cohesive modules then receive more reports than those in less cohesive ones do but also send fewer. There may still be plenty of top-down communication, with agents in more cohesive modules reporting to others in less cohesive ones, but in aggregate communication will be bottom up.

DEFINITION. *Suppose there are communication links from agents in module $\mathcal{M}_m$ to agents in module $\mathcal{M}_{m'}$ but not the reverse. Then communication from module $\mathcal{M}_m$ to $\mathcal{M}_{m'}$ is "top down" if $x_m > x_{m'}$ and "bottom up" if $x_{m'} > x_m$. Communication is "bottom up in aggregate" if there are more pairs of modules that engage in bottom-up than top-down communication.*

We can then state the result.

PROPOSITION 3. *If the optimal sender and receiver hierarchies are the reverse of each other, and the receiver ranking is strict, communication is bottom up in aggregate.*

In the proof of the proposition we further show that communication is *strictly* bottom up—that there are strictly more pairs of modules that engage in bottom-up than top-down communication—unless communication costs are so low that some agents in each module tell their states to all the others.

Suppose next that the optimal sender and receiver hierarchies are identical, that is, for all $i, j \in \mathcal{N}$, agent i outranks agent j in the receiver hierarchy if and only if he also outranks him in the sender hierarchy. Once again speaking loosely, agents in more cohesive modules then do not only receive more reports than those in less cohesive ones do, they also send more. As a result, communication is no longer necessarily bottom up in aggregate. Instead, the agents in the most cohesive modules now form a core whose members engage in intense communication with each other, while agents in the less cohesive modules form a periphery whose members communicate only sparsely with the core, and even less with each other.

DEFINITION. *A communication network has a "core-periphery structure" if the set of modules can be partitioned into a core and a periphery such that (i.) an agent in the core tells his state to all the other agents in the core and possibly to agents in the periphery, (ii.) an agent in the periphery either does not tell his state to all the agents in the core, or is not told all their states, or both, and (iii.) an agent in the periphery does not tell his state to other agents in the periphery.*

We then have the following result.

PROPOSITION 4. *If the optimal sender and receiver hierarchies are identical, the communication network has a core-periphery structure in which the core consists of the most cohesive modules.*

As we noted in the introduction, core-periphery structures are common in social and communication networks. We defer a real-world illustration to the end of the next section, where we will return to these sort of structures.

5 Application—Mirroring Hypothesis

The result that the optimal organization of modular production is hierarchical contrasts with the Mirroring Hypothesis. As we discussed earlier, the Mirroring Hypothesis conjectures that the optimal way to organize modular production is to mirror the production function, to enable intense communication within modules and accept sparse communication across. In our setting, an organization mirrors its production function if the principal places communication links within modules but not across.

DEFINITION. *An organization "mirrors" the production function if agent* $i \in \mathcal{N}$ *tells agent* $j \in \mathcal{N}$ *about his state if and only if they belong to the same module.*

The Boeing Company's experience with the 787 Dreamliner illustrates the Mirroring Hypothesis and why it may not always hold.[5] The Dreamliner was designed to be modular precisely because it allowed Boeing to outsource the development and production of most modules to independent suppliers, many of which were scattered around the globe (see Figure 5). Suppliers delivered the finished modules to Boeing's factory in Everett, where its workers put them together with the tail fin, the only major module still made by Boeing itself. To the extent that firm- and country boundaries hamper communication, this way of organizing the production of the Dreamliner is broadly in line with the Mirroring Hypothesis.

The intention of Boeing's organizational strategy was to speed up the development of the Dreamliner and save production costs. This is not what happened. As an article in Reuters reported at the time (Peterson 2011): "*On a blustery and drizzly December afternoon in the Pacific Northwest, about 20 airplanes sat engineless and inert near the runway at a Boeing manufacturing plant...The program that produced these unfinished 787s is nearly three years behind schedule and, by some estimates, at least several billion dollars over budget.*"[6] The underlying reason for these delays and cost

[5] This account is based on Peterson (2011) and Brown and Garthwaite (2016). See also Tadelis and Williamson (2013).

[6] In line with the description above, the article goes on to say: "*The 787 is not merely a historic feat of engineering. The program also marks Boeing's departure from its own time-honored manufacturing practices. Instead of drawing*

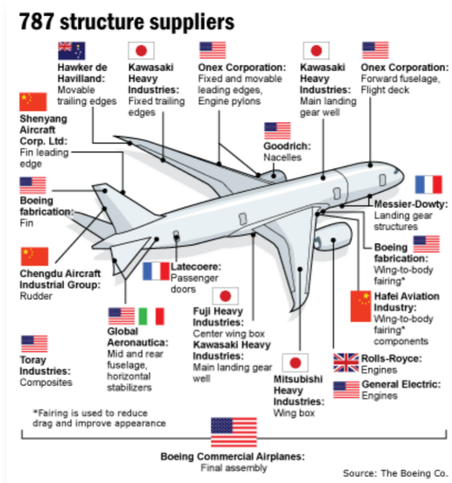

Figure 5: Suppliers of Boeing's 787 Dreamliner (replication of Figure 1 in Koster et al. (2011)).

overruns were coordination problems among the suppliers and between them and Boeing. These problems proved so severe that Boeing was eventually forced to abandon its organizational strategy and bring the production of some modules back in-house: "*Some of the parts arriving in Everett did not fit together, and late deliveries by producers of crucial sections of the plane stopped the entire assembly process…As a result, Boeing was forced to reverse some of its original outsourcing decisions; for example, in 2009 it spent \$1 billion in cash and credit to acquire its fuselage manufacturing partner Vought Aircraft Industries*" (Brown and Garthwaite 2016, p.12). Even when products are highly modular, therefore, mirroring might fail because the need to coordinate across modules necessitates intense communication between agents working on different ones.

In our setting, mirroring is a corner solution of the principal's problem in which no agent tells his state to any agent in another module. The force that pushes optimal communication away from this corner solution is the supermodularity of the coordination multipliers and thus the agents' expected revenues. Across-module communication always improves decision-making but it does so especially given the intensity of within-module communication that arises under mirroring. For across-module communication to be unprofitable nevertheless, the degree of coupling has to be sufficiently low. Just how low it needs to be depends on the characteristics of the production network, as described in the next proposition.

---

primarily from its traditional pool of aircraft engineers, mechanics and laborers that runs generations deep in the Puget Sound region around Seattle, Boeing leads an international team of suppliers and engineers from the United States, Japan, Italy, Australia, France and elsewhere, who make components that Boeing workers in the United States put together."

PROPOSITION 5. *Mirroring is optimal if and only if* $t \le \min_{i \in \mathcal{N}} \bar{t}_i$, *where* $\bar{t}_i > 0$ *is the threshold degree of coupling above which it is optimal for agent* i *to tell his state to agents in modules other than his own and below which it is not. Adding modules to the production function decreases the threshold* $\bar{t}_i$, *as does increasing the module characteristics* $n_{m'}$ *or* $p_{m'}$ *for any* $m' \neq m(i)$.

For the Mirroring Hypothesis to hold, then, there cannot be too many modules, and none of the modules can consist of too many decisions or require too much coordination, or else some agents should tell their states to agents in other modules. Arguably, this is why mirroring failed at Boeing. Its production function was modular but still very complex, comprising many modules, some of which involved many decisions that required a high degree of coordination. Our model suggests that, in such a firm, across-module communication can be essential, even when the degree of coupling is low.

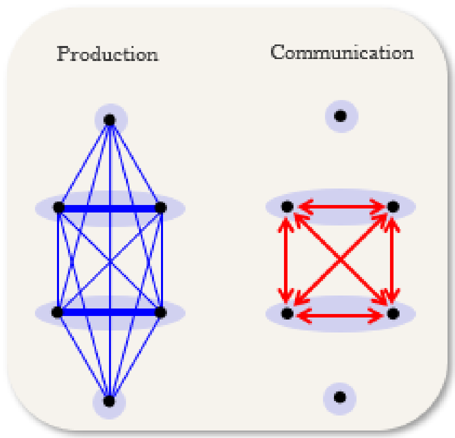

Figure 6: Illustration of an optimal communication network that partially mirrors the production network. Left panel—A production network consisting of four modules, two with two decisions and two with only one (where the shaded areas highlight the modules). Right panel—The optimal communication network. Parameter values: need for coordination within the two-decision modules is 0.5, degree of coupling is 0.01, and the value of autonomous adaptation is 1 for all nodes; communication costs can take any value $\gamma \in (0.000434, 0.000801)$.

A broader notion of the Mirroring Hypothesis allows for *modular-like* organizations, ones that contain clusters of modules whose agents communicate with each other but not with agents outside of the cluster. The management literature refers to such arrangements as "*partial mirroring.*"

DEFINITION. *An organization "partially mirrors" the production function if the set of modules can be partitioned into subsets such that (i.) agent* $i \in \mathcal{N}$ *tells agent* $j \in \mathcal{N}$, *who belongs to a* *different module, about his state if and only if their modules belong to the same subset and (ii.) there is at least one "cluster," that is, a subset that contains two or more modules.*

Partial mirroring can be optimal in our setting, and Figure 6 provides an example. If it is, though, it has to take a particular form.

PROPOSITION 6. *When partial mirroring is optimal, the organization contains only one cluster of modules, and the modules that form the cluster are the most cohesive ones.*

This result follows from the optimality of hierarchies. If there were multiple clusters, communication would not be hierarchical, which cannot be optimal. Suppose, for instance, that one cluster consists of modules $\mathcal{M}_1$ and $\mathcal{M}_2$ and another of modules $\mathcal{M}_3$ and $\mathcal{M}_4$. If it is optimal for an agent in $\mathcal{M}_1$ to tell his state to agents in $\mathcal{M}_2$ but not to those in $\mathcal{M}_3$, then $\mathcal{M}_2$ has to be more cohesive than $\mathcal{M}_3$. But if $\mathcal{M}_2$ is more cohesive than $\mathcal{M}_3$, it cannot be optimal for an agent in $\mathcal{M}_4$ to tell his state to agents in $\mathcal{M}_3$ but not to those in $\mathcal{M}_2$. In contrast, the existence of a single cluster is consistent with the optimal design of communication networks, provided it consists of the most cohesive modules. The optimal communication network then has a core-periphery structure of the type we discussed at the end of the previous section, albeit one with no communication between the core and the periphery.

This takes us back to Boeing and its decision to respond to the failure of its initial organizational strategy by insourcing the production of some modules, such as the fuselage, while continuing to leave the production of others to its suppliers. Provided, again, that information flows more freely within firms than across, this response created a core-periphery structure in which the in-house modules formed the core and the outsourced ones the periphery. To the extent that the tail fin and the fuselage, as well as the other modules Boeing brought in-house, were the most cohesive ones, this response is consistent with the optimal design of communication networks in our model.

6 Robustness

Some of the assumptions in the model are critical for our results, while others are merely convenient. The assumption that communication costs are lower within modules than across is one of the convenient ones and allows us to streamline the exposition. Beyond convenience, this assumption captures the notion that, because of physical proximity and shared expertise, agents working on the same module may find it easier to explain their local conditions to each other than to those working on other modules. The characterization result extends readily to an alternative specification in which the costs of a communication link are the same within and across modules. We examine this alternative in Appendix B.

Another assumption that is not essential for our results is that communication is perfectly informative, which is an assumption we share with Calvó-Armengol and de Martí Beltran (2009). Suppose, instead, that when agent i tells agent j his state, agent j learns θ_i with probability $q \in (0,1]$ and pure noise otherwise. To keep higher-order beliefs simple, suppose further that both parties know whether communication was effective. Finally, suppose that if agent i tells his state to multiple agents, they either all learn his state or none of them do. The effectiveness of communication, in other words, is specific to the sender and does not vary across receivers. In Appendix C we show that allowing for imperfect communication in this manner boils down to re-scaling the communication costs: the principal's problem is still given by (7) with communication costs γ replaced by γ/q.

Among the more critical assumptions, it is useful to distinguish between those that matter for the separability result in Proposition 1 and those that matter for the characterization in Proposition 2. An assumption that is essential for the separability result is that states are independent, as in Calvó-Armengol, de Martí, and Prat (2015). If states were correlated, an agent who is told about one state would also learn some information about the other states, reducing the benefit of telling him about them. As a result, the problems of who should be told about each state would be interdependent, breaking the separability result.

Another assumption that is critical for the separability result is that agents do not re-transmit information, which is an assumption we share with Calvó-Armengol and de Martí Beltran (2009), Calvó-Armengol, de Martí, and Prat (2015), and Herskovic and Ramos (2020). If agent i talks to another agent, he tells him about his state θ_i but not about any other information he may have been told about. This assumption captures the notion that, even though we model each agent's state as simply a number, it refers to a complex set of conditions and circumstances that only the associated agent can describe appropriately. If agents were able to re-transmit information, the separability result would fail. Communication links from agent i to other agents would then affect the overall cost, and thus optimal placement, of communication links from any other agent j.

Finally, the separability result depends on the absence of incentive conflicts. As mentioned earlier, this is the defining assumption of the literature on team theory. To see how incentive conflicts affect the separability result, suppose that each agent only internalizes a fraction $\mu \in [0,1]$ of the needs to coordinate and acts as if the production network were given by $\mu \boldsymbol{P}$ rather than $\boldsymbol{P}$ (for instance because they put more weight on their own revenue or profits, as in Athey and Roberts (2001) and Alonso, Dessein, and Matouschek (2008)). The rest of the model is as in Section 3. The next proposition shows how such incentive conflicts affect the separability result.

PROPOSITION 7. *If agents internalize only a fraction $\mu \in [0,1]$ of the needs to coordinate, an optimal communication network solves*

$$\max_{\boldsymbol{C}} \sum_{i=1}^{N} a_i \mathrm{Cov}\left(d_i^*, \theta_i\right) + (1-\mu) \sum_{i=1}^{N} \sum_{j=1}^{N} p_{ij} \mathrm{Cov}(d_i^*, d_j^*) - \gamma \sum_{i=1}^{N} \sum_{j=1}^{N} m_{ij} c_{ij}, \tag{9}$$

where

$$\mathrm{Cov}\left(d_i^*, \theta_i\right) = a_i \sigma_i^2 \omega_{ii}\left(\boldsymbol{C}_i, \mu\right)$$

and

$$\mathrm{Cov}(d_i^*, d_j^*) = \sum_{s=1}^{N} a_s^2 \sigma_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right),$$

and where $\omega_{ij}\left(\boldsymbol{C}_j, \mu\right)$ denotes the ijth entry of $\left(\boldsymbol{I} - \left(\mathrm{diag}\boldsymbol{C}_j\right) \mu \boldsymbol{P} \left(\mathrm{diag}\boldsymbol{C}_j\right)\right)^{-1}$ and m_{ij} is a dummy variable that is equal to one if and only if agents i and j belong to different modules.

The only new term in the principal's objective function (9) is the weighted sum of the covariances between each decision pair. Its presence implies that if agents are biased against coordination, it is no longer enough for the principal to ensure that each decision is sufficiently adapted to its state. Instead, she also needs to take into account how communication affects coordination and what she can do to ensure decisions co-vary more strongly with each other. The challenge this property poses is that the extent to which two decisions co-vary with each other depends which states both decisions makers are informed about. The principal can, therefore, no longer consider each agent in isolation and ask whom he should tell about his state. She has to consider all agents at once and take into account how communication links from one agent affect the optimal location of such links from the others. Since the objective function continues to be supermodular, the principal can still use standard algorithms to solve for optimal communication networks in polynomial time. Finding an analytical solution, however, becomes more challenging.

The key assumption in the entire paper is that the production function has a non-overlapping community structure. This assumption allows us to capture the notion that products are modular, which are the type of products we are interested in. To generalize this structure, one could allow for different degrees of coupling—different ts—for decisions in different pairs of modules. Since a module may consist of a single decision, though, such a production function would constitute a general, unweighted network with little structure to base a characterization of optimal communication on. To see what can still be said in this case, suppose that the production network $\boldsymbol{P}$ can take any form, provided it still satisfies $p_{ii} = 0$, $p_{ij} = p_{ji}$, and $\sum_{j=1}^{N} p_{ij} < 1$ for all $i, j \in \mathcal{N}$. The separability result in Proposition 1, and the lemmas that precede it, continue to hold for these more general production functions. As such, the principal can still determine the optimal communication network by considering each agent in isolation. Moreover, the principal's objective is still supermodular and can, therefore, be maximized using standard algorithms.

What can no longer hold is the characterization of optimal communication networks in Proposition 2, which is specific to the non-overlapping community structure. Optimal communication can now take many forms and need not give rise to hierarchies. The specific form it takes depends on the specific structure of the production network. There are, however, some properties of optimal communication networks that hold across production networks.

PROPOSITION 8. *As long as the production network* $\boldsymbol{P}$ *satisfies* $p_{ii} = 0$, $p_{ij} = p_{ji}$, *and* $\sum_{j=1}^{N} p_{ij} < 1$, *optimal communication networks* $\boldsymbol{C}^*$ *are increasing in the value of autonomous adaptation* $a_i^2\sigma_i^2$ *and the needs for coordination* p_{ij} *for all* $i, j \in \mathcal{N}$, *and decreasing in communication costs* γ.

The proposition shows that, in general, the principal will only ever respond to an increase in the value of adaptation or the need for coordination, or a decrease in the cost of communication, by adding communication links. These comparative statics hold because the principal's objective function is supermodular and has either increasing or decreasing differences in the various parameters (Topkis 1978, Milgrom and Shannon 1994).

7 Extension—Nested Modules

We started this paper with Herbert Simon's observation that complex systems tend to be made up of modules. As we noted in the introduction, Simon did not stop there and, instead, went on to argue that such systems often have a nearly decomposable structure, in which each module is itself made up of sub-modules, each sub-module is made up of sub-sub-modules, and so on, with interactions being stronger within any type of module than across. To conclude the formal analysis, we adapt our model to take a first step towards exploring optimal communication when production networks have such a nested, modular structure. Our goal is to demonstrate that our model can be adapted to explore related issues without providing a comprehensive exploration of nested, modular networks. As such, we focus on a symmetric version of the type of system Simon described.

Specifically, suppose the total number of decisions, and associated states and agents, is given by $N = n^k$, where n and k are positive integers. Each decision belongs to one level-1 module, which consists of n decisions, each level-1 module belongs to one level-2 module, which consists of n level-1 modules, and so on. The highest level module is the level-k module, which consists of n

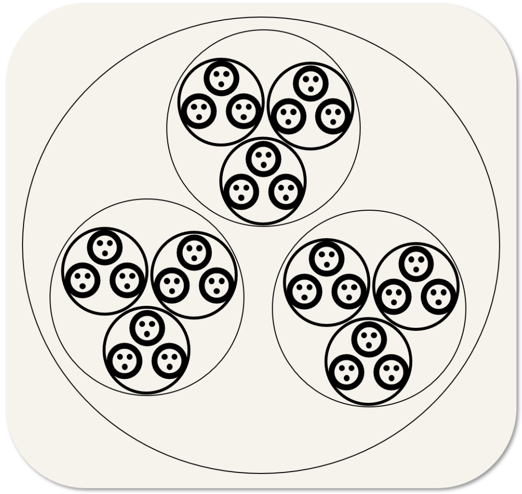

Figure 7: A production network with a nested, modular structure.

level-k–1 modules. The need for coordination $p_{ij} > 0$ between any decisions d_i and $d_j \neq d_i$ is given by t_1 if they belong to the same level-1 module, by $t_2 < t_1$ if they belong to the same level-2 module but not the same level-1 module, and so on. As in the main model, we set p_{ii} equal to zero for all $i \in \mathcal{N}$ and require $\sum_{j=1}^{N} p_{ij} < 1$. For simplicity, suppose that the value of autonomous adaptation $a_i^2 \sigma_i^2$ is the same for all $i \in \mathcal{N}$ and that the cost of any agent telling his state to another is the same $\gamma > 0$, irrespective of the modules they belong to. The rest of the model is the same as the main model.

We illustrate the production network of this extension in Figure 7, where each node represents a decision, and associated state and agent, and circles indicate the different types of modules. Nodes belonging to the same, thickest circle form a level-1 module, those belonging to the same, second-thickest circle form a level-2 module, and so on.

The separability result continues to hold in this setting, allowing us to again solve the principal's problem by determining whom each agent should tell about his state. Moreover, since all agents now enter the production network in the same way, it suffices to ask whom any one agent should tell about his state. The next proposition gives the answer.

PROPOSITION 9. *If the production network is symmetric and has a nested, modular structure, optimal communication networks are characterized by a threshold $\kappa \in \{1, 2, ..., k\}$. Agent $i \in \mathcal{N}$ tells his state to agent $j \in \mathcal{N}$ if and only if they belong to the same level-κ module.*

The proposition shows that optimal communication again has a threshold structure. Agents tell their states only to their closest neighbors, those with whom they have the highest needs for coordination. Such communication is reminiscent of the type of partial mirroring we discussed above, with intense communication taking place within each level-κ module but not across. We leave an exploration of the extent to which optimal communication deviates from this simple form once one allows for asymmetries for future research.

8 Conclusions

The rise of modular production over the last 60 years has been widely observed and documented and has been explored extensively in management and computer science. The goal of this paper was to take a first step towards understanding the economic implications of the rise of modular production.

As a first step, we focused on the immediate implications of modular production for the internal organization of firms and abstracted from broader implications for their boundaries and the structure of industries. Even in this narrow context, many open questions remain. An important practical issue we put aside is the role of *interfaces* which ensure that different modules fit with each other. One way to think about such interfaces in our model is as a limited set of decisions that are made and announced before agents make the remaining ones.

Another widely-discussed issue we did not address is *parallel processing*, the notion that modular production allows firms to accelerate production by having different agents work on different modules simultaneously (Parnas 1972). One way to get at this issue in our model is to suppose that the principal can hire agents and decide which decisions each agent is in charge of. Each agent first spends time learning the states associated with his decisions, taking one period per state to do so. After all the agents have learned their states, they make their decisions simultaneously and without spending any further time on communication. A patient principal would hire a single agent and have him make the first-best decisions after N periods but an impatient principal may prefer to hire M agents, put each in charge of one module, and have them make worse decisions sooner. We leave the investigation of both interfaces and parallel processing, as well as other issues related to internal organization, for future research.

The impact of modular production on the economy is unlikely to be confined to changes in the internal organization of firms. Baldwin and Clark (1997), for instance, observe that, while the introduction of the System/360 did lead to immediate changes in IBM's internal organization, its more enduring impact was to cause entry into the computer industry in the following decades. The entrants were often small, entrepreneurial firms that focused on the development and production of individual modules and whose innovative products allowed them to compete successfully with

IBM's own, in-house module makers. In this telling, the introduction of the System/360 in the 1960s sowed the seeds for the subsequent disintegration of IBM and the other large mainframe manufacturers and gave rise to the competitive and innovative computer industry of today.[7] There are many reasons why modular production may affect the boundaries of firms and the structure and inventiveness of industries and we leave their exploration for future research.

A question that goes beyond the impact of modular production on economic activity is what explains its rise in the first place. Simon argued that modularity facilitates adaptation by confining adaptive changes to individual modules within a system. The argument that modularity allows parallel processing provides another reason why it may have adaptive advantages. In line with these intuitions, firms such as IBM explain their development of modular products with the need to adapt quickly to the changing capabilities of their suppliers and needs of their customers. Yet, a full explanation for the rise of modular production also needs to account for its costs. It may be easier to adapt a modular product to its environment but, for a given environment, one would expect limitations in across-module interactions to affect its performance. After all, products have not always been modular, and even today many are not, suggesting that such designs also have significant downsides. Answering the questions of when and why firms develop modular products, and what trade-offs they face when they are doing so, would require moving beyond one of the foundational economic modeling assumptions, that production functions are given by nature and not designed by firms. As such, it is the most challenging question this paper highlights and, like the other open questions we sketched above, we leave it for future research.

Finally, this paper also raises empirical issues. Newly available data sets contain detailed information about communication between employees within firms (Impink, Prat, and Sadun (2021) and Yang et al. (2021)). Our model makes specific predictions about the pattern of such communication in firms that make modular products. In particular, the prediction that optimal communication has a hierarchical structure has a number of implications that are, at least in principle, observable, such as the emergence of core-periphery structures and the absence of multiple cores or clusters. While testing the model is naturally difficult, we hope that this paper provides additional stimulation and direction to the emerging empirical literature on within-firm communication.

[7] As Baldwin and Clark (1997) observe: "*But modularity also undermined IBM's dominance in the long run, as new companies produced their own so-called plug-compatible modules—printers, terminals, memory, software, and eventually even the central processing units themselves—that were compatible with, and could plug right into, the IBM machines. By following IBM's design rules but specializing in a particular area, an upstart company could often produce a module that was better than the ones IBM was making internally. Ultimately, the dynamic, innovative industry that has grown up around these modules developed entirely new kinds of computer systems that have taken away most of the mainframe's market share.*"

References

[1] Alexander, Christopher. 1964. *Notes on the Synthesis of Form.* Harvard University Press.

[2] Alonso, Ricardo, Wouter Dessein, and Niko Matouschek. 2008. When Does Coordination Require Centralization? *American Economic Review.* 98(1): 145–79.

[3] Arrow, Kenneth. 1974. *The Limits of Organization.* W.W. Norton & Company.

[4] Athey, Susan and John Roberts. 2001. Organizational Design: Decision Rights and Incentive Contracts. *American Economic Review.* 91(2): 200–5.

[5] Baccara, Mariagiovanna and Heski Bar-Isaac. 2008. How to Organize Crime. *The Review of Economic Studies.* 75(4): 1039–67.

[6] Bala, Venkatesh and Sanjeev Goyal. 2000. A Noncooperative Model of Network Formation. *Econometrica.* 68(5): 1181–229.

[7] Baldwin, Carliss and Kim Clark. 1997. Managing in an Age of Modularity. *Harvard Business Review.* 75(5): 84–93.

[8] ——— and ——— . 2000. *Design Rules: The Power of Modularity* (Vol.1). MIT Press.

[9] Ballester, Coralio, Antoni Calvó-Armengol, and Yves Zenou. 2006. Who's Who in Networks. Wanted: The Key Player. *Econometrica.* 74(5): 1403–17.

[10] Bergemann, Dirk, Tibor Heumann, and Stephen Morris. 2017. Information and Interaction. Cowles Foundation Discussion Paper No. 2088.

[11] Borgatti, Stephen and Martin Everett. 2000. Models of Core/Periphery Structures. *Social Networks.* 21(4): 375–95.

[12] Brown, Jennifer and Craig Garthwaite. 2016. Global Aircraft Manufacturing, 2002–2011. *Kellogg School of Management Case Study.* KEL 938.

[13] Calvó-Armengol, Antoni and Joan de Martí Beltran. 2009. Information Gathering in Organizations: Equilibrium, Welfare, and Optimal Network Structure. *Journal of the European Economic Association.* 7(1): 116–61.

[14] ———, Joan de Martí, and Andrea Prat. 2015. Communication and Influence. *Theoretical Economics.* 10(2): 649–90.

[15] Colfer, Lyra and Carliss Baldwin. 2016. The Mirroring Hypothesis: Theory, Evidence, and Exceptions. *Industrial and Corporate Change*. 25(5): 709–38.

[16] Conway, Melvin. 1968. How Do Committees Invent? *Datamation*. 14(4): 84–93.

[17] Dessein, Wouter and Tano Santos. 2006. Adaptive Organizations. *Journal of Political Economy*. 114(5): 956–95.

[18] ———, Andrea Galeotti, and Tano Santos. 2016. Rational Inattention and Organizational Focus. *American Economic Review*. 106(6): 1522–36.

[19] Fortunato, Santo. 2010. Community Detection in Graphs. *Physics Reports*. 486(3–5): 75–174.

[20] Garicano, Luis and Andrea Prat. 2013. Organizational Economics with Cognitive Costs. *Advances in Economics and Econometrics*. 1: 342–88.

[21] ——— and Timothy Van Zandt. 2013. Hierarchies and the Division of Labor. In *The Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts. Princeton University Press: 604–54.

[22] Garud, Raghu, Arun Kumaraswamy, and Richard Langlois, eds. 2009. *Managing in the Modular Age: Architectures, Networks, and Organizations*. John Wiley & Sons.

[23] Golub, Benjamin and Stephen Morris. 2017. Expectations, Networks, and Conventions. Mimeo.

[24] Goyal, Sanjeev and Adrien Vigier. 2014. Attack, Defence, and Contagion in Networks. *The Review of Economic Studies*. 81(4): 1518–42.

[25] Guimera, Roger, Stefano Mossa, Adrian Turtschi, and Luis Nunes Amaral. 2005. The Worldwide Air Transportation Network: Anomalous Centrality, Community Structure, and Cities' Global Roles. *Proceedings of the National Academy of Sciences*. 102(22): 7794–99.

[26] Henderson, Rebecca and Kim Clark. 1990. Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. *Administrative Science Quarterly*. 35(1): 9–30.

[27] Herskovic, Bernard and Joao Ramos. 2020. Acquiring Information through Peers. *American Economic Review*. 110(7): 2128–52.

[28] Impink, Stephen, Andrea Prat, and Raffaella Sadun. 2021. Communication within Firms: Evidence from CEO Turnovers. NBER Working Paper Series, No. 29042.

[29] Jackson, Matthew and Asher Wolinsky. 1996. A Strategic Model of Social and Economic Networks. *Journal of Economic Theory.* 71(1): 44–74.

[30] Koster, Jean, Ewald Kraemer, Claus-Dieter Munz, Dries Verstraete, K. C. Wong, and Alec Velazco. 2011. Workforce Development for Global Aircraft Design. *ASME International Mechanical Engineering Congress and Exposition*, 54877: 661–71.

[31] Langlois, Richard and Paul Robertson. 1992. Networks and Innovation in a Modular System: Lessons from the Microcomputer and Stereo Component Industries. *Research Policy.* 21(4): 297–313.

[32] Marschak, Jacob and Roy Radner. 1972. *Economic Theory of Teams.* Yale University Press.

[33] McCord, Kent and Steven Eppinger. 1993. Managing the Integration Problem in Concurrent Engineering. MIT Sloan School of Management, Cambridge, MA, Working Paper No. 3594.

[34] Meunier, David, Renaud Lambiotte, Alex Fornito, Karen Ersche, and Edward Bullmore. 2009. Hierarchical Modularity in Human Brain Functional Networks. *Frontiers in Neuroinformatics.* 3: 37.

[35] Milgrom, Paul and Chris Shannon. 1994. Monotone Comparative Statics. *Econometrica.* 62(1): 157–80.

[36] Morris, Stephen. 2000. Contagion. *The Review of Economic Studies.* 67(1): 57–78.

[37] Murota, Kazuo. 2003. *Discrete Convex Analysis.* Society for Industrial and Applied Mathematics.

[38] Parnas, David. 1972. On the Criteria to be Used in Decomposing Systems in Modules. In *Pioneers and Their Contributions to Software Engineering*, eds. Manfred Broy and Ernst Denert. Springer: 479–98.

[39] Peterson, Kyle. 2011. Special Report: A Wing and a Prayer: Outsourcing at Boeing. *Reuters.* January 20, 2011.

[40] Rombach, Puck, Mason Porter, James Fowler, and Peter Mucha. 2017. Core-Periphery Structure in Networks (Revisited). *SIAM Journal on Applied Mathematics*. 59(3): 619–46.

[41] Sanchez, Ron and Joseph Mahoney. 1996. Modularity, Flexibility, and Knowledge Management in Product and Organization Design. *Strategic Management Journal*. 17(2): 63–76.

[42] Simon, Herbert. 1962. The Architecture of Complexity. *Proceedings of the American Philosophical Society*. 106 (6): 467–82.

[43] Tadelis, Steven and Oliver Williamson. 2013. Transaction Cost Economics. In *The Handbook of Organizational Economics*, eds. Robert Gibbons and John Roberts. Princeton University Press: 159–90.

[44] Talamàs, Eduard and Omer Tamuz. 2017. Network Cycles and Welfare. Mimeo.

[45] Thompson, James. 1967. Organizations in Action: Social Science Bases of Administrative Theory. McGraw-Hill.

[46] Topkis, Donald. 1978. Minimizing a Submodular Function on a Lattice. *Operations Research*. 26(2): 305–21.

[47] Yang, Longqi, David Holtz, Sonia Jaffe, Siddharth Suri, Shilpi Sinha, Jeffrey Weston, Connor Joyce, Neha Shah, Kevin Sherman, Brent Hecht, and Jaime Teevan. 2021. The Effects of Remote Work on Collaboration Among Information Workers. *Nature Human Behaviour*: 1–12.

Appendix A: Proofs

We first introduce some notation. Throughout the appendix, for notational compactness, we will denote by $\mathrm{E}_i[\cdot]$ the expectation over $\theta = (\theta_1, \ldots, \theta_N)$ given the information agent i has under communication network $\boldsymbol{C}$. That is, for a random variable Z, $\mathrm{E}_i[Z] \equiv \mathrm{E}\left[Z \middle| \boldsymbol{C}_{(i)}\right]$. Next, a strategy for agent i is a mapping $\tilde{d}_i : [-D, D]^N \to [-D, D]$, where $\tilde{d}_i(\theta)$ denotes the decision that agent i makes in state θ. We denote a strategy profile by $\tilde{d} = \times_{i=1}^N \tilde{d}_i$.

LEMMA 1. *Equilibrium decisions are unique and given by*

$$d_i^* = \sum_{j=1}^{N} a_j \omega_{ij} (\boldsymbol{C}_j) \theta_j \textit{ for all } i \in \mathcal{N},$$

where $\omega_{ij}(\boldsymbol{C}_j)$ *denotes the* ijth *entry of* $(\boldsymbol{I} - (\operatorname{diag} \boldsymbol{C}_j) \boldsymbol{P} (\operatorname{diag} \boldsymbol{C}_j))^{-1}$.

Proof of Lemma 1. This proof parallels the approach of Golub and Morris (2017), Appendix A1: We take the communication network $\boldsymbol{C}$ as given and show that $d^* = \times_{i=1}^N d_i^*$ is the unique strategy profile that survives iterated elimination of strictly dominated strategies and is therefore the unique Bayesian-Nash equilibrium.

Step 1: *Show that there is a unique Bayesian-Nash equilibrium by showing that there is a unique strategy profile that survives iterated elimination of strictly dominated strategies.*

Given any $\boldsymbol{C}$, the game played by the agents is a game of strategic complements: if we denote

$$\hat{d}_i\left(\theta, \tilde{d}_{-i}\right) = a_i \theta_i + \sum_{j=1}^{N} p_{ij} \mathrm{E}_i\left[\tilde{d}_j(\theta)\right]$$

agent i's best response to the strategy profile $\tilde{d}_{-i}$ in state θ, then $\hat{d}_i$ is increasing in each $\tilde{d}_j$ under the partial order given by $\tilde{d}_j \succeq \tilde{d}'_j$ if and only if $\tilde{d}_j(\theta) \geq \tilde{d}'_j(\theta)$ for all θ.

Define the set $S_i(k)$ to be the set of i's pure strategies surviving k rounds of iterated elimination of strictly dominated strategies. By assumption, $d_i(\theta) \in [-D, D]$, so the first set in the sequence is

$$S_i(0) = \left\{ \tilde{d}_i \middle| -D \leq \tilde{d}_i(\theta) \leq D \text{ for all } \theta \right\}.$$

Next, as this is a game of strategic complements, an upper bound on $S_i(1)$ is i's best response to the maximal strategy profile $\tilde{d}_{-i} \in S_{-i}(0)$, where $\tilde{d}_{-i} = \times_{j \neq i} \tilde{d}_j$ and $S_{-i}(k) = \prod_{j \neq i} S_j(k)$, and a lower bound on $S_i(1)$ is i's best response to the minimal strategy profile $\tilde{d}_{-i} \in S_{-i}(0)$. That is,

$$S_i(1) = \left\{ \tilde{d}_i \middle| a_i \theta_i - \sum_{j=1}^{N} p_{ij} D \leq \tilde{d}_i(\theta) \leq a_i \theta_i + \sum_{j=1}^{N} p_{ij} D \right\}.$$

Next, suppose that for $k > 1$, the set $S_i(k)$ takes the form

$$S_i(k) = \left\{ \tilde{d}_i \middle| \underline{d}_i^k(\theta) \leq \tilde{d}_i(\theta) \leq \overline{d}_i^k(\theta) \text{ for all } \theta \right\},$$

where

$$\begin{aligned}
\overline{d}_i^k(\theta) &= a_i\theta_i + \sum_{m=1}^{k-1} \beta_{im} + \sum_{j_1=1}^{N} \sum_{j_2=1}^{N} \cdots \sum_{j_k=1}^{N} p_{ij_1} p_{j_1 j_2} \cdots p_{j_{k-1} j_k} D \\
\underline{d}_i^k(\theta) &= a_i\theta_i + \sum_{m=1}^{k-1} \beta_{im} - \sum_{j_1=1}^{N} \sum_{j_2=1}^{N} \cdots \sum_{j_k=1}^{N} p_{ij_1} p_{j_1 j_2} \cdots p_{j_{k-1} j_k} D,
\end{aligned}$$

and

$$\beta_{im} = \sum_{j_1=1}^{N} \cdots \sum_{j_m=1}^{N} p_{ij_1} p_{j_1 j_2} \cdots p_{j_{m-1} j_m} a_{j_m} E_i E_{j_1} \cdots E_{j_{m-1}} [\theta_{j_m}].$$

Then an upper bound on $S_i(k+1)$ is agent i's best response to the maximal strategy profile $\tilde{d}_{-i} \in S_{-i}(k)$, and a lower bound on $S_i(k+1)$ is agent i's best response to the minimal strategy profile $\tilde{d}_{-i} \in S_{-i}(k)$. That is,

$$S_i(k+1) = \left\{ \tilde{d}_i \middle| \underline{d}_i^{k+1}(\theta) \leq \tilde{d}_i(\theta) \leq \overline{d}_i^{k+1}(\theta) \text{ for all } \theta \right\}.$$

To show that the upper and lower bounds of $S_i(k)$ converge to the same value, we show that

$$\lim_{k\to\infty} \sum_{j_1=1}^{N} \sum_{j_2=1}^{N} \cdots \sum_{j_k=1}^{N} p_{ij_1} p_{j_1 j_2} \cdots p_{j_{k-1} j_k} D = 0.$$

This term converges to zero as long as the row sum of the production matrix to the kth power, $\boldsymbol{P}^k$, converges to zero as $k \to \infty$. This result follows since $\sum_{j=1}^{N} p_{ij} < 1$ for all i, and therefore the spectral radius of $\boldsymbol{P}$ is strictly less than one. By the sandwich theorem, we therefore have

$$\lim_{k\to\infty} \underline{d}_i^k(\theta) = \lim_{k\to\infty} \overline{d}_i^k(\theta) = a_i\theta_i + \sum_{m=1}^{\infty} \beta_{im}.$$

This result implies that $\lim_{k\to\infty} S_i(k)$ is a singleton for all i. As this is a supermodular game, the resulting strategy profile is the unique Bayesian-Nash equilibrium of the game.

Step 2: *Show that the unique Bayesian-Nash equilibrium strategy profile is a linear combination of $\theta_1, \ldots, \theta_N$, that is, $d_i^*(\theta) = \sum_{j=1}^{N} \alpha_{ij}\theta_j$ for some scalars $\{\alpha_{ij}\}_{j=1}^{N}$.*

First, observe that $\mathrm{E}_i \mathrm{E}_{j_1} \cdots \mathrm{E}_{j_{m-1}} [\theta_{j_m}]$ is zero if some $j \in \{i, j_1, \ldots, j_{m-1}\}$ does not know θ_{j_m} under $\boldsymbol{C}$, and $\mathrm{E}_i \mathrm{E}_{j_1} \cdots \mathrm{E}_{j_{m-1}} [\theta_{j_m}] = \theta_{j_m}$ if all $j \in \{i, j_1, \ldots, j_{m-1}\}$ know θ_{j_m} under $\boldsymbol{C}$. This result follows by an induction argument and the law of iterated expectations. For the $m = 1$

case, $\mathrm{E}_i[\theta_{j_1}] = 0$ if i does not know θ_{j_1}, and $\mathrm{E}_i[\theta_{j_1}] = \theta_{j_1}$ if i does know θ_{j_1}. Next, suppose all $j \in \{i, j_1, \ldots, j_{m-1}\}$ know $\theta_{j_{m+1}}$. Then $\mathrm{E}_{j_m}\mathrm{E}_i\mathrm{E}_{j_1}\cdots\mathrm{E}_{j_{m-1}}\left[\theta_{j_{m+1}}\right] = \mathrm{E}_{j_m}\left[\theta_{j_{m+1}}\right]$, which is 0 if j_m does not know $\theta_{j_{m+1}}$ and is $\theta_{j_{m+1}}$ if all $\{i, j_1, \ldots, j_m\}$ know $\theta_{j_{m+1}}$. Finally, suppose there is some $j \in \{i, j_1, \ldots, j_{m-1}\}$ who does not know $\theta_{j_{m+1}}$. Then $\mathrm{E}_{j_m}\mathrm{E}_i\mathrm{E}_{j_1}\cdots\mathrm{E}_{j_{m-1}}\left[\theta_{j_{m+1}}\right] = \mathrm{E}_{j_m}[0] = 0$.

The result in the previous paragraph ensures that each β_{im} from Step 1 is a linear combination of $\theta_1, \ldots, \theta_N$, and therefore $d_i^*(\theta) = a_i\theta_i + \sum_{m=1}^{\infty}\beta_{im} = \sum_{j=1}^{N}\alpha_{ij}\theta_j$ for some scalars $\{\alpha_{ij}\}_{j=1}^{N}$.

Step 3: *Show that* $\alpha_{ij} = a_j\omega_{ij}(\boldsymbol{C}_j)\theta_j$, where $\omega_{ij}(\boldsymbol{C}_j)$ *denotes the* ijth *entry of the matrix* $(\boldsymbol{I} - (\operatorname{diag}\boldsymbol{C}_j)\boldsymbol{P}(\operatorname{diag}\boldsymbol{C}_j))^{-1}$.

The network associated with $(\operatorname{diag}\boldsymbol{C}_j)\boldsymbol{P}(\operatorname{diag}\boldsymbol{C}_j)$ is the subgraph of the production network induced by nodes that know θ_j, and therefore $\omega_{ij}(\boldsymbol{C}_j)$ is the sum of the values of all walks from node i to node j on the production network that pass only through nodes that know θ_j.

Note that $p_{ij_1}p_{j_1j_2}\cdots p_{j_{m-1}j_m}$ describes the value of a walk of length m from node i to node j_m on the production network. If any node in a walk $ij_1, j_1j_2, \cdots, j_{m-1}j_m$ does not know θ_{j_m}, then from the argument in step 2, $\mathrm{E}_i\mathrm{E}_{j_1}\cdots\mathrm{E}_{j_{m-1}}[\theta_{j_m}] = 0$. Otherwise, $\mathrm{E}_i\mathrm{E}_{j_1}\cdots\mathrm{E}_{j_{m-1}}[\theta_{j_m}] = \theta_{j_m}$. Thus, β_{im} is the sum of the values of all walks of length m from node i to node j_m on the production network that pass only through nodes that know θ_{j_m}. The result then follows. ■

COROLLARY 1. *The weight* $a_i\omega_{ii}(\boldsymbol{C}_i)$ *that agent* i*'s decision* d_i^* *puts on his state* θ_i *satisfies* $\omega_{ii}(\boldsymbol{I}_i)a_i = a_i$, *where* I_i *is the* i*th row of an* $N \times N$ *identity matrix, and is increasing and supermodular in* $\boldsymbol{C}_i$.

Proof of Corollary 1. This result follows from the proofs of Lemma 1 and Proposition 8. ■

LEMMA 2. *Under equilibrium decision-making, expected revenue is given by*

$$R(\boldsymbol{C}) \equiv \mathrm{E}\left[r\left(d_1^*, \ldots, d_N^*\right)\right] = \sum_{i=1}^{N} a_i \mathrm{Cov}\left(d_i^*, \theta_i\right),$$

where $\mathrm{Cov}(d_i^*, \theta_i) = a_i\sigma_i^2\omega_{ii}(\boldsymbol{C}_i)$.

Proof of Lemma 2. Given optimal decision making, revenue in state θ can be written as

$$\sum_{i=1}^{N} a_i d_i^* \theta_i - \sum_{i=1}^{N} d_i^* \left[d_i^* - a_i\theta_i - \sum_{j=1}^{N} p_{ij} d_j^*\right].$$

Next, substitute in the best responses $d_i^* = a_i\theta_i + \sum_{j=1}^{N} p_{ij}\mathrm{E}_i\left[d_j^*\right]$. The term in square brackets is therefore equal to $\sum_{j=1}^{N} p_{ij}\left[\mathrm{E}_i\left[d_j^*\right] - d_j^*\right]$, and revenue can be written as

$$\sum_{i=1}^{N} a_i d_i^* \theta_i - \sum_{i=1}^{N}\sum_{j=1}^{N} p_{ij} d_i^* \left[\mathrm{E}_i\left[d_j^*\right] - d_j^*\right].$$

We will now show that agent i's optimal decision d_i^* is orthogonal to the error with which he predicts agent j's optimal decision, $\mathrm{E}_i\left[d_j^*\right]-d_j^*$, and therefore the second set of terms is zero in expectation. To see why this result is true, take expectations of the second term, and plug in $d_i^*=\sum_{s=1}^N a_s\omega_{is}\left(\boldsymbol{C}_s\right)\theta_s$ and $\mathrm{E}\left[d_j^*\right]=\sum_{t=1}^N a_t\omega_{jt}\left(\boldsymbol{C}_t\right)\mathrm{E}_i\left[\theta_t\right]$. We therefore have

$$\begin{aligned}
& \mathrm{E}\left[\sum_{i=1}^N\sum_{j=1}^N p_{ij}\sum_{s=1}^N a_s\omega_{is}\left(\boldsymbol{C}_s\right)\theta_s\left[\sum_{t=1}^N a_t\omega_{jt}\left(\boldsymbol{C}_t\right)\mathrm{E}_i\left[\theta_t\right]-\sum_{t=1}^N a_t\omega_{jt}\left(\boldsymbol{C}_t\right)\theta_t\right]\right] \\
= & \sum_{i=1}^N\sum_{j=1}^N p_{ij}\sum_{s=1}^N \omega_{is}\left(\boldsymbol{C}_s\right)a_s\sum_{t=1}^N \omega_{jt}\left(\boldsymbol{C}_t\right)a_t\mathrm{E}\left[\theta_s\mathrm{E}_i\left[\theta_t\right]-\theta_s\theta_t\right] \\
= & \sum_{i=1}^N\sum_{j=1}^N p_{ij}\sum_{s=1}^N \omega_{is}\left(\boldsymbol{C}_s\right)\omega_{js}\left(\boldsymbol{C}_s\right)a_s^2\mathrm{E}\left[\theta_s\mathrm{E}_i\left[\theta_s\right]-\theta_s\theta_t\right].
\end{aligned}$$

In the last equality, the cross terms cancel by independence. The binary information structure ensures that each term in the final line is zero, since either agent i knows θ_s, in which case the final term in square brackets is zero, or agent i does not know θ_s, in which case $\omega_{is}\left(\boldsymbol{C}_s\right)=0$.

Expected revenue is therefore

$$\mathrm{E}\left[\sum_{i=1}^N a_i d_i^*\theta_i\right]=\sum_{i=1}^N a_i\mathrm{Cov}\left(d_i^*,\theta_i\right)=\sum_{i=1}^N a_i^2\sigma_i^2\omega_{ii}\left(\boldsymbol{C}_i\right),$$

since

$$\mathrm{Cov}\left(d_i^*,\theta_i\right)=\mathrm{E}\left[d_i^*\theta_i\right]-\mathrm{E}\left[d_i^*\right]\mathrm{E}\left[\theta_i\right]=\mathrm{E}\left[d_i^*\theta_i\right]=\mathrm{E}\left[\sum_{j=1}^N a_j\omega_{ij}\left(\boldsymbol{C}_j\right)\theta_j\theta_i\right]=a_i\omega_{ii}\left(\boldsymbol{C}_i\right)\mathrm{E}\left[\theta_i^2\right],$$

where the final equality follows by independence of the states. ■

PROPOSITION 1. *An optimal communication network solves the principal's problem* (2) *if and only if it solves the N independent subproblems*

$$\max_{\boldsymbol{C}_i} R_i\left(\boldsymbol{C}_i\right)-\gamma\sum_{j=1}^N m_{ij}c_{ij} \text{ subject to } c_{ii}=1 \text{ for all } i\in\mathcal{N},$$

where m_{ij} is a dummy variable equal to one if and only if agents i and $j\neq i$ belong to different modules.

Proof of Proposition 1. Optimal communication networks maximize expected revenues minus communication costs. Using the expected revenue expression derived in Lemma 2, optimal communication networks solve

$$\max_{\boldsymbol{C}}\left[\sum_{i=1}^N a_i^2\sigma_i^2\omega_{ii}\left(\boldsymbol{C}_i\right)-\gamma\sum_{i=1}^N\sum_{j=1}^N m_{ij}c_{ij}\right].$$

Since $\omega_{ii}(\boldsymbol{C}_i)$ depends only on $\boldsymbol{C}_i$ and not the rest of the communication network $\boldsymbol{C}$, and the objective is additively separable in i, a communication network $\boldsymbol{C}^*$ solves this problem if and only if $\boldsymbol{C}_i^*$ solves

$$\max_{\boldsymbol{C}_i} a_i^2 \sigma_i^2 \omega_{ii}(\boldsymbol{C}_i) - \gamma \sum_{j=1}^{N} m_{ij} c_{ij}$$

for all $i \in \mathcal{N}$. ■

LEMMA 3. *Suppose agent 1 tells his state to all agents in modules $\mathcal{M}_1, \ldots, \mathcal{M}_\ell$, for $\ell \in \{1, \ldots, M\}$ and to none of the agents in other modules. Agent 1's expected revenue is then given by*

$$R_1(\boldsymbol{C}_1(\ell)) = a_1^2 \sigma_1^2 \left(\frac{1-(n_1-2)p_1}{(1+p_1)(1-(n_1-1)p_1)} + \frac{t^2 x_1^2 \left(\sum_{m=1}^{\ell} n_m x_m - n_1 x_1\right)}{(1-tn_1 x_1)\left(1-t\sum_{m=1}^{\ell} n_m x_m\right)} \right),$$

where

$$x_m \equiv \frac{1}{1-(n_m-1)p_m + n_m t} \text{ for } m = 1, \ldots, M.$$

Proof of Lemma 3. Denote the modules whose agents know state θ_1 by $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_\ell$, where $\mathcal{M}_1$ is agent 1's own module. From Lemma 2, the expected revenue generated by agent 1 is $R_1(\boldsymbol{C}_1(\ell)) = a_1^2\sigma_1^2\omega_{11}(\boldsymbol{C}_1(\ell))$. We will derive $\omega_{11}(\boldsymbol{C}_1(\ell))$ in four steps.

Step 1: *Derive a representation of $\omega_{11}(\boldsymbol{C}_1(\ell))$ as the value of walks on a modified module-level production network, and show that $\omega_{11}(\boldsymbol{C}_1(\ell)) = \det \boldsymbol{V}(\ell) / \det \boldsymbol{Q}(\ell)$ for some matrices $\boldsymbol{V}(\ell)$ and $\boldsymbol{Q}(\ell)$.*

The value $\omega_{11}(\boldsymbol{C}_1(\ell))$ is the sum of the values of all walks from node 1 back to itself on the subgraph of the production network consisting of nodes in modules whose agents know state θ_1. Denote this value by v. Next, let v_k be the sum of the values of all walks from a node in module k to node 1 on this same subgraph. These values can be written recursively as a system of equations.

$$\begin{aligned} v &= 1 + p_1(n_1-1)v_1 + tn_2v_2 + \cdots + tn_\ell v_\ell \\ v_1 &= p_1 v + p_1(n_1-2)v_1 + tn_2v_2 + \cdots + tn_\ell v_\ell \\ &\vdots \\ v_\ell &= tv + t(n_1-1)v_1 + tn_2v_2 + \cdots + p_\ell(n_\ell-1)v_\ell. \end{aligned}$$

The right-hand side of the first equation describes the value of all walks from node 1 in the following way: the first term, 1, is the value of walks that pass only through node 1; the second term, $p_1(n_1-1)v_1$, is the value of all walks that initially pass to one of the n_1-1 other nodes in module $\mathcal{M}_1$; the $k+1$th term, $tn_k v_k$, is the value of all walks that initially pass to one of the n_k nodes in module $\mathcal{M}_k$. The right-hand side of the second equation captures the value of all walks from a node $j \neq 1$ in module $\mathcal{M}_1$ back to node 1 in the following way: the first term is the value of walks that initially pass back to node 1; the second term is the value of all walks that initially pass to one of the other $n_1 - 2$ nodes in module $\mathcal{M}_1$; the $k+1$th term is the value of all walks that initially pass to one of the n_k nodes in module $\mathcal{M}_k$. The remaining equations are interpreted analogously.

This system of $\ell + 1$ equations can be written in matrix form:

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & -p_1(n_1-1) & -tn_2 & \cdots & -tn_\ell \\ -p_1 & 1-p_1(n_1-2) & -tn_2 & \cdots & -tn_\ell \\ -t & -t(n_1-1) & 1-p_2(n_2-1) & \cdots & -tn_\ell \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -t & -t(n_1-1) & -tn_2 & \cdots & 1-p_\ell(n_\ell-1) \end{bmatrix} \begin{bmatrix} v \\ v_1 \\ v_2 \\ \vdots \\ v_\ell \end{bmatrix}.$$

Denote this $(\ell+1) \times (\ell+1)$ matrix by $\boldsymbol{Q}(\ell)$. Then v is the $(1,1)$ element of the inverse matrix $\boldsymbol{Q}(\ell)^{-1}$, and by the definition of a matrix inverse, $v = \det \boldsymbol{V}(\ell) / \det \boldsymbol{Q}(\ell)$, where $\boldsymbol{V}(\ell)$ is the matrix obtained by removing the first row and column of $\boldsymbol{Q}(\ell)$.

Step 2: *Show that* $\det \boldsymbol{Q}(\ell) = \frac{\det \boldsymbol{A}}{x_2 \cdots x_\ell}\left(1 - z\sum_{j=2}^{\ell} n_j x_j\right)$, *where* $\det \boldsymbol{A} = (1+p_1)(1-p_1(n_1-1))$, *and* $z = t\left(1 + t\frac{(1+p_1)n_1}{\det \boldsymbol{A}}\right)$.

We use the following result for the determinant of a block matrix,

$$\det \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix} = \det(\boldsymbol{A}) \det\left(\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}\right).$$

Partition matrix $\boldsymbol{Q}(\ell)$ so that $\boldsymbol{A}$ is defined as $\begin{bmatrix} 1 & -p_1(n_1-1) \\ -p_1 & 1-p_1(n_1-2) \end{bmatrix}$, and $\boldsymbol{B}$, $\boldsymbol{C}$, and $\boldsymbol{D}$ are defined accordingly. Then

$$\begin{aligned} \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B} &= \begin{bmatrix} -t & -t(n_1-1) \\ -t & -t(n_1-1) \\ \vdots & \vdots \\ -t & -t(n_1-1) \end{bmatrix} \begin{bmatrix} 1 & -p_1(n_1-1) \\ -p_1 & 1-p_1(n_1-2) \end{bmatrix}^{-1} \begin{bmatrix} -tn_2 & -tn_3 & \cdots & -tn_\ell \\ -tn_2 & -tn_3 & \cdots & -tn_\ell \end{bmatrix} \\ &= \frac{t^2(1+p_1)n_1}{\det \boldsymbol{A}} \begin{bmatrix} n_2 & n_3 & \cdots & n_\ell \\ n_2 & n_3 & \cdots & n_\ell \\ \vdots & \vdots & \ddots & \vdots \\ n_2 & n_3 & \cdots & n_\ell \end{bmatrix}. \end{aligned}$$

Define $z = t\left(1 + t\frac{(1+p_1)n_1}{\det \boldsymbol{A}}\right)$, and recall from the statement of the lemma that $x_s^{-1} = 1 - p_s(n_s - 1) + tn_s$. Using these expressions, we can write

$$\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B} = \begin{bmatrix} x_2^{-1} - zn_2 & -zn_3 & \cdots & -zn_\ell \\ -zn_2 & x_3^{-1} - zn_3 & \cdots & -zn_\ell \\ \vdots & \vdots & \ddots & \vdots \\ -zn_2 & -zn_2 & \cdots & x_\ell^{-1} - zn_\ell \end{bmatrix}.$$

For $\ell = 2$, the determinant of this matrix is $x_2^{-1} - zn_2$. For $\ell \geq 3$, since the determinant is invariant to subtracting one row from another, we can subtract row 2 from row 1, row 3 from row 2, row 4 from row 3, and so on, to rewrite $\det\left(\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}\right)$ as

$$\det \begin{bmatrix} x_2^{-1} & -x_3^{-1} & 0 & \cdots & 0 \\ 0 & x_3^{-1} & -x_4^{-1} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -x_\ell^{-1} \\ -zn_2 & -zn_3 & -zn_4 & \cdots & x_\ell - zn_\ell \end{bmatrix}.$$

Denote this matrix by $\boldsymbol{Z}_{2\to\ell}$. Denote by $\boldsymbol{Z}_{3\to\ell}$ the matrix derived by deleting the first row and column from $\boldsymbol{Z}_{2\to\ell}$ and denote $\boldsymbol{Z}_{k+1\to\ell}$ the matrix derived by deleting the first row and column from $\boldsymbol{Z}_{k\to\ell}$.

Using the cofactor expansion of the determinant along the first column, we can rewrite the expression for $\det \boldsymbol{Z}_{2\to\ell}$ as

$$\det \boldsymbol{Z}_{2\to\ell} = x_2^{-1} \det \boldsymbol{Z}_{3\to\ell} - zn_2 x_3^{-1} x_4^{-1} \cdots x_\ell^{-1}.$$

For the second term, we use the result that the determinant of a lower-triangular matrix is the product of its trace. More generally,

$$\det \boldsymbol{Z}_{k\to\ell} = x_k^{-1} \det \boldsymbol{Z}_{k+1\to\ell} - zn_k x_{k+1}^{-1} x_{k+2}^{-1} \cdots x_\ell^{-1}.$$

Note that $\det \boldsymbol{Z}_{\ell-1\to\ell} = \frac{1}{x_{\ell-1}x_\ell}\left(1 - z\left(n_{\ell-1}x_{\ell-1} + n_\ell x_\ell\right)\right)$, so iteratively plugging in the expression for $\det \boldsymbol{Z}_{\ell-k\to\ell}$ into the expression for $\det \boldsymbol{Z}_{\ell-k-1\to\ell}$ gives us

$$\det \boldsymbol{Z}_{2\to\ell} = \frac{1}{x_2 \cdots x_\ell}\left(1 - z\sum_{j=2}^{\ell} n_j x_j\right),$$

and since $\det \boldsymbol{Q}(\ell) = (\det \boldsymbol{A}) \det\left(\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}\right) = (\det \boldsymbol{A})(\det \boldsymbol{Z}_{2\to\ell})$, we have

$$\det \boldsymbol{Q}(\ell) = \frac{\det \boldsymbol{A}}{x_2 \cdots x_\ell}\left(1 - z\sum_{j=2}^{\ell} n_j x_j\right),$$

which completes this step.

Step 3: *Show that* $\det \boldsymbol{V}(\ell) = \frac{\det \tilde{\boldsymbol{A}}}{x_2 \cdots x_\ell}\left(1 - \tilde{z}\sum_{j=2}^{\ell} n_j x_j\right)$, *where* $\det \tilde{\boldsymbol{A}} = 1 - p_1(n_1 - 2)$, *and* $\tilde{z} = t\left(1 + t\frac{n_1 - 1}{\det \tilde{\boldsymbol{A}}}\right)$.

This step proceeds similarly to step 2. Recall that $\boldsymbol{V}(\ell)$ is the matrix derived by eliminating the first row and column from the matrix $\boldsymbol{Q}(\ell)$. Partition $\boldsymbol{V}(\ell)$ into the block matrix $\begin{bmatrix} \tilde{\boldsymbol{A}} & \tilde{\boldsymbol{B}} \\ \tilde{\boldsymbol{C}} & \tilde{\boldsymbol{D}} \end{bmatrix}$ by letting $\tilde{\boldsymbol{A}} = 1 - p_1(n_1 - 2)$ and setting $\tilde{\boldsymbol{B}}$, $\tilde{\boldsymbol{C}}$, and $\tilde{\boldsymbol{D}}$ accordingly. Then

$$\tilde{\boldsymbol{C}}\tilde{\boldsymbol{A}}^{-1}\tilde{\boldsymbol{B}} = \frac{t^2(n_1 - 1)}{\det \tilde{\boldsymbol{A}}}\begin{bmatrix} n_2 & n_3 & \cdots & n_\ell \\ n_2 & n_3 & \cdots & n_\ell \\ \vdots & \vdots & \ddots & \vdots \\ n_2 & n_3 & \cdots & n_\ell \end{bmatrix}.$$

Define $\tilde{z} = t\left(1 + t\frac{n_1 - 1}{\det \tilde{\boldsymbol{A}}}\right)$. Then

$$\tilde{\boldsymbol{D}} - \tilde{\boldsymbol{C}}\tilde{\boldsymbol{A}}^{-1}\tilde{\boldsymbol{B}} = \begin{bmatrix} x_2^{-1} - \tilde{z}n_2 & -\tilde{z}n_3 & \cdots & -\tilde{z}n_\ell \\ -\tilde{z}n_2 & x_3^{-1} - \tilde{z}n_3 & \cdots & -\tilde{z}n_\ell \\ \vdots & \vdots & \ddots & \vdots \\ -\tilde{z}n_2 & -\tilde{z}n_2 & \cdots & x_\ell^{-1} - \tilde{z}n_\ell \end{bmatrix},$$

which is the same expression as for $\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}$ in step 2, except that we have replaced z with $\tilde{z}$. Applying the same argument as in step 2, we therefore have

$$\det \boldsymbol{V}(\ell) = \det\left(\tilde{\boldsymbol{A}}\right)\det\left(\tilde{\boldsymbol{D}} - \tilde{\boldsymbol{C}}\tilde{\boldsymbol{A}}^{-1}\tilde{\boldsymbol{B}}\right) = \frac{\det\left(\tilde{\boldsymbol{A}}\right)}{x_2 \cdots x_\ell}\left(1 - \tilde{z}\sum_{j=2}^{\ell} n_j x_j\right),$$

which completes this step.

Step 4: *Show that* $\omega_{11}(\boldsymbol{C}_1(\ell)) = \frac{1-(n_1-2)p_1}{(1+p_1)(1-(n_1-1)p_1)} + \frac{t^2 x_1^2\left(\sum_{m=1}^{\ell} n_m x_m - n_1 x_1\right)}{(1-t n_1 x_1)\left(1 - t\sum_{m=1}^{\ell} n_m x_m\right)}$.

We will decompose $\omega_{11}(\boldsymbol{C}_1(\ell))$ into two terms: (*i.*) the value of all walks from node 1 back to node 1 that pass only through nodes in module $\mathcal{M}_1$ plus (*ii.*) the value of all walks from node 1 back to node 1 that pass through nodes in modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$. If we define $\boldsymbol{Q}(1) = \begin{bmatrix} 1 & -p_1(n_1 - 1) \\ -p_1 & 1 - p_1(n_1 - 2) \end{bmatrix}$, then the value of (*i.*) is $\frac{\det \boldsymbol{V}(1)}{\det \boldsymbol{Q}(1)}$. This decomposition allows us to rewrite $\omega_{11}(\boldsymbol{C}_1(\ell))$ as

$$\omega_{11}(\boldsymbol{C}_1(\ell)) = \frac{\det \boldsymbol{V}(\ell)}{\det \boldsymbol{Q}(\ell)} = \frac{\det \boldsymbol{V}(1)}{\det \boldsymbol{Q}(1)} + \frac{\det \boldsymbol{V}(\ell)\det \boldsymbol{Q}(1) - \det \boldsymbol{V}(1)\det \boldsymbol{Q}(\ell)}{\det \boldsymbol{Q}(\ell)\det \boldsymbol{Q}(1)}.$$

The first expression is

$$\frac{1-p_1(n_1-2)}{(1+p_1)(1-p_1(n_1-1))},$$

and the second expression is

$$\begin{aligned} & \frac{t^2\sum_{j=2}^{\ell} n_j x_j}{(1-p_1(n_1-1))(1-p_1(n_1-1))\left(1-t\left(1+t\frac{n_1}{1-p_1(n_1-1)}\right)\sum_{j=2}^{\ell} n_j x_j\right)} \\ = \quad & \frac{t^2 x_1^2 \sum_{j=2}^{\ell} n_j x_j}{(1-tn_1x_1)\left(1-t\sum_{j=1}^{\ell} n_j x_j\right)}, \end{aligned}$$

where the first term follows from substitution and rearrangement, and the second term uses the the definition of x_1 to rewrite $1-p_1(n_1-1)=\frac{1}{x_1}(1-tn_1x_1)$. These two expressions complete step 4. The lemma then follows because $R_1(\boldsymbol{C}_1(\ell))=a_1^2\sigma_1^2\omega_{11}(\boldsymbol{C}_1(\ell))$. ■

PROPOSITION 2. *Optimal communication is characterized by N thresholds $\lambda_i \geq 0$, one for each agent $i \in N$. Agent i tells agent j about his state if and only if they belong to the same module, or the cohesion of agent j's module is above agent i's threshold, that is, $x_{m(j)} \geq \lambda_i$. The threshold λ_i is increasing in marginal communication costs and decreasing in the value of autonomous adaptation $a_i^2\sigma_i^2$, the need to coordinate the decisions with agent i's module $p_{m(i)}$, and the size of his module $n_{m(i)}$.*

Proof of Proposition 2. To establish the threshold result, we will argue that if agent i informs a module $\mathcal{M}_j$ with cohesion x_j but not a module $\mathcal{M}_{j'}$ with cohesion $x_{j'}$, then it must be the case that $x_j > x_{j'}$. To do so, we will show that the expected per-node incremental revenue of informing both modules $\mathcal{M}_j$ and $\mathcal{M}_{j'}$ is greater than the expected per-node incremental revenue of informing only module $\mathcal{M}_{j'}$ whenever $x_j > x_{j'}$. This result will imply that whenever it is optimal to inform module $\mathcal{M}_{j'}$, it is also optimal to inform module $\mathcal{M}_j$.

To make this argument, we use the notation from the proof of Lemma 3. The expected per-node incremental revenue of informing one more module, $\mathcal{M}_{\ell+1}$, about state θ_1, given that modules $\mathcal{M}_1,\ldots,\mathcal{M}_\ell$ are informed, is

$$\frac{1}{n_{\ell+1}}a_1^2\sigma_1^2\left[\frac{\det \boldsymbol{V}(\ell+1)}{\det \boldsymbol{Q}(\ell+1)}-\frac{\det \boldsymbol{V}(\ell)}{\det \boldsymbol{Q}(\ell)}\right]=\frac{1}{n_{\ell+1}}a_1^2\sigma_1^2\frac{\det \boldsymbol{V}(\ell+1)\det \boldsymbol{Q}(\ell)-\det \boldsymbol{V}(\ell)\det \boldsymbol{Q}(\ell+1)}{\det \boldsymbol{Q}(\ell+1)\det \boldsymbol{Q}(\ell)}.$$

Using the expressions for $\det \boldsymbol{Q}(\cdot)$ and $\det \boldsymbol{V}(\cdot)$ derived in Lemma 3, we can rewrite this expression:

$$a_1^2\sigma_1^2\frac{t^2x_1^2x_{\ell+1}}{\left(1-t\sum_{j=1}^{\ell} n_j x_j\right)\left(1-t\sum_{j=1}^{\ell+1} n_j x_j\right)}.$$

Analogously, the expected per-node incremental revenue of informing modules $\mathcal{M}_{\ell+1}$ and $\mathcal{M}_{\ell+2}$ about state θ_1, given that modules $\mathcal{M}_1, \ldots, \mathcal{M}_\ell$ are informed, is

$$\frac{1}{n_{\ell+1}+n_{\ell+2}} \frac{a_1^2 \sigma_1^2 t^2 x_1^2 \left(n_{\ell+1} x_{\ell+1}+n_{\ell+2} x_{\ell+2}\right)}{\left(1-t \sum_{j=1}^{\ell} n_j x_j\right)\left(1-t \sum_{j=1}^{\ell+2} n_j x_j\right)}.$$

Then the expected per-node incremental revenue of additionally informing modules $\mathcal{M}_{\ell+1}$ and $\mathcal{M}_{\ell+2}$ minus the expected per-node incremental revenue of additionally informing only module $\mathcal{M}_{\ell+1}$ about θ_1 can be written as follows.

$$\begin{aligned}
& \frac{1}{n_{\ell+1}+n_{\ell+2}} \frac{a_1^2 \sigma_1^2 t^2 x_1^2 \left(n_{\ell+1} x_{\ell+1}+n_{\ell+2} x_{\ell+2}\right)}{\left(1-t \sum_{j=1}^{\ell} n_j x_j\right)\left(1-t \sum_{j=1}^{\ell+2} n_j x_j\right)}-a_1^2 \sigma_1^2 \frac{t^2 x_1^2 x_{\ell+1}}{\left(1-t \sum_{j=1}^{\ell} n_j x_j\right)\left(1-t \sum_{j=1}^{\ell+1} n_j x_j\right)} \\
= \quad & \frac{a_1^2 \sigma_1^2 x_1^2 t^2 n_{\ell+2}\left[\left(x_{\ell+2}-x_{\ell+1}\right)\left(1-t \sum_{j=1}^{\ell+1} n_j x_j\right)+t\left(n_{\ell+1}+n_{\ell+2}\right) x_{\ell+1} x_{\ell+2}\right]}{\left(n_{\ell+1}+n_{\ell+2}\right)\left(1-t \sum_{j=1}^{\ell} n_j x_j\right)\left(1-t \sum_{j=1}^{\ell+1} n_j x_j\right)\left(1-t \sum_{j=1}^{\ell+2} n_j x_j\right)}.
\end{aligned}$$

This expression is necessarily positive if $x_{\ell+2} \geq x_{\ell+1}$. Notice that this condition is independent of which modules (besides $\mathcal{M}_1$) already know state θ_1.

We can rank the modules according to their values of cohesion x_s and show, by contradiction, that if it is optimal to inform any module with cohesion x_j, then it is optimal to inform all modules with $x_s \geq x_j$. Suppose it is optimal to inform some subset of modules, $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_\ell$, but there is some $x_K \geq x_i$ for some $K \notin \{2, \ldots, \ell\}$ and some $i \in \{2, \ldots, \ell\}$. If all modules $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_\ell$ but not $\mathcal{M}_i$ know state θ_1, the above inequality tells us that expected per-node incremental revenue is higher if both modules $\mathcal{M}_i$ and $\mathcal{M}_K$ also know θ_1 than if only $\mathcal{M}_i$ knows θ_1. This contradicts the optimality of informing modules $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_\ell$.

The threshold λ_1 is given by the value at which the expected revenue of informing all modules $\mathcal{M}_s$ with $x_s \geq \lambda_1$ minus the cost of informing those modules is highest. If profits are negative from informing any subset of modules, then $\lambda_1 > x_s$ for all $s \in \{2, \ldots, M\}$. Recall that the expected revenue generated by agent 1 when modules $\mathcal{M}_1, \ldots, \mathcal{M}_\ell$ know state θ_1 is

$$a_1^2 \sigma_1^2 \left(\frac{1-p_1\left(n_1-2\right)}{\left(1+p_1\right)\left(1-p_1\left(n_1-1\right)\right)}+\frac{t^2 x_1^2 \sum_{j=2}^{\ell} n_j x_j}{\left(1-t n_1 x_1\right)\left(1-t \sum_{j=1}^{\ell} n_j x_j\right)}\right),$$

where $x_s = \frac{1}{1-p_s(n_s-1)+t n_s}$. The first term in brackets is increasing in n_1 and p_1, and the second term is increasing in x_1 and n_1, where x_1 is increasing in p_1 and n_1, since $p_s \geq t$ by assumption. Thus, the expected revenue from informing modules $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_\ell$ is increasing in a_1, σ_1^2, n_1, and p_1. Similarly, the expected per-node incremental revenue from additionally informing modules $\mathcal{M}_{\ell+1}, \ldots, \mathcal{M}_{\ell+k}$ given modules $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_\ell$ are already informed is increasing in a_1, σ_1^2, n_1, and p_1. Following an increase in a_1, σ_1^2, n_1, and p_1, since the incremental revenue from additionally informing modules is increasing in a_1, σ_1^2, n_1, and p_1, then the profit from informing all modules with $x_s \geq \lambda_1$ continues to be higher than the profit from informing a subset of those modules. ■

COROLLARY 2. *Optimal communication gives rise to a receiver hierarchy among agents. For any agents $i, j, k \in \mathcal{N}$ who belong to different modules, if $x_{m(i)} \geq x_{m(j)}$, then agent j is told about agent k's state only if agent i also is.*

Proof of Corollary 2. This result follows directly from Proposition 2. ■

COROLLARY 3. *Optimal communication gives rise to a sender hierarchy among agents. For any agents $i, j, k \in \mathcal{N}$ who belong to different modules, if $\lambda_i \leq \lambda_j$, then agent j tells agent k about his state only if agent i also does.*

Proof of Corollary 3. This result follows directly from Proposition 3. ■

PROPOSITION 3. *If the optimal sender and receiver hierarchies are the reverse of each other, and the receiver ranking is strict, communication is bottom up in aggregate.*

Proof of Proposition 3. Label the modules by their cohesion, with the most-cohesive module labeled $\mathcal{M}_1$ and the least-cohesive module labeled $\mathcal{M}_M$: $x_1 > x_2 > \cdots > x_M$. Let us first introduce some definitions. We define the *receiver ranking* $\succeq_R$ by $i \succeq_R j$ if and only if $x_{m(i)} \geq x_{m(j)}$ and the *strict receiver ranking* $i \succ_R j$ if and only if $x_{m(i)} > x_{m(j)}$. We say that the *sender ranking* $\succeq_S$ is the *reverse of the receiver ranking* if $i \succeq_R j$ if and only if $j \succeq_S i$. Because agents within a module can have different sender rankings, the property that the sender ranking is the reverse of the receiver ranking implies that $i \succ_R j \Rightarrow j \succ_S i$, but we can have that $i \succ_S j$ and $i \sim_R j$. Finally, we say that a communication link from i to j is *unilateral top down* if $x_{m(i)} > x_{m(j)}$, and there is no communication link from j to i. We establish the result in three steps. The first step establishes an intermediate result, and the last two steps apply it to two cases, establishing the proposition.

Step 1: *Show that communication from i to j can be unilateral top down only if (i.) agents i and j are in adjacent modules (i.e., $m(j) = m(i) + 1$) and (ii.) no agent in module $\mathcal{M}_{m(i)}$ informs any agent in any module $\mathcal{M}_{m(i)+s}$ for $s \geq 2$.*

To establish the first part of the claim, let us suppose communication from i to j is unilateral top down, and $m(j) > m(i) + 1$. Then because everyone in module j has the same receiver rank, it must be the case that i informs all the agents in module $m(j)$, and because $x_{m(j)} < x_{m(j)-1}$, it must also be the case that i informs all the agents in module $m(j) - 1$. Since $i \succ_R j$, it must be the case that $j \succ_S i$, so j must inform all agents in module $m(j) - 1$ as well as any agent k with $x_{m(k)} > x_{m(j)-1}$, which includes any agent in module $m(i)$. Agent j must therefore inform agent i, which is a contradiction. If communication from i to j is unilateral top down, it must be the case that $m(j) = m(i) + 1$.

For the second part of the claim, suppose communication from i to j is unilateral top down with $m(j) = m(i) + 1$, and suppose some agent in module $\mathcal{M}_{m(i)}$ informs some agent k with $m(k) > m(j)$. Then agent j must also inform agent k and any agent in a more cohesive module. Module $\mathcal{M}_{m(i)}$ is more cohesive than module $\mathcal{M}_{m(k)}$. Agent j must therefore inform agent i, which is a contradiction. If communication from i to j is unilateral top down, it must be the case that no agent in module $\mathcal{M}_{m(i)}$ informs any agent in any module $\mathcal{M}_{m(i)+s}$ for $s \geq 2$.

Step 2: *Show that if there is top-down communication from module $\mathcal{M}_m$, $m \leq M-2$, to module $\mathcal{M}_{m+1}$, there is bottom-up communication from module $\mathcal{M}_{m+2}$ to module $\mathcal{M}_{m+1}$ and from module $\mathcal{M}_{m+2}$ to module $\mathcal{M}_m$.*

Suppose there is top-down communication from some module $\mathcal{M}_m$ to the module $\mathcal{M}_{m+1}$. Then by the second part of Step 1, there is no agent in module $\mathcal{M}_m$ who informs any agent in module $\mathcal{M}_{m+2}$. Since $x_m > x_{m+2}$, agents in module $\mathcal{M}_{m+2}$ outrank agents in module $\mathcal{M}_m$ in the sender ranking and therefore must also inform some agents in module $\mathcal{M}_{m+1}$. Since no agent in module $\mathcal{M}_{m+1}$ informs any agent in module $\mathcal{M}_m$, they must also not inform any agent in module $\mathcal{M}_{m+2}$, so there is bottom-up communication from module $\mathcal{M}_{m+2}$ to module $\mathcal{M}_{m+1}$.

Moreover, since agents in module $\mathcal{M}_{m+2}$ inform agents in module $\mathcal{M}_{m+1}$, they must also inform agents with a higher receiver rank. That is, they must also inform agents in module $\mathcal{M}_m$, and so there is also bottom-up communication from module $\mathcal{M}_{m+2}$ to module $\mathcal{M}_m$. Thus, if there is top-down communication from a module $\mathcal{M}_m$ to module $\mathcal{M}_{m+1}$, there is bottom-up communication from module $\mathcal{M}_{m+2}$ to modules $\mathcal{M}_{m+1}$ and from module $\mathcal{M}_{m+2}$ to module $\mathcal{M}_m$.

Step 3: *Show that if there is top-down communication from module $\mathcal{M}_{M-1}$ to module $\mathcal{M}_M$, there is is bottom-up communication from module $\mathcal{M}_{M-1}$ to module $\mathcal{M}_{M-2}$ and from module $\mathcal{M}_M$ to module $\mathcal{M}_{M-2}$.*

Suppose there is top-down communication from module $\mathcal{M}_{M-1}$ to module $\mathcal{M}_M$. Given that no agent in module $\mathcal{M}_M$ informs any agent in module $\mathcal{M}_{M-1}$, and module $\mathcal{M}_M$ dominates module $\mathcal{M}_{M-2}$ in the sender ranking, then no agent in module $\mathcal{M}_{M-2}$ informs any agent in module $\mathcal{M}_{M-1}$ or $\mathcal{M}_M$. By the receiver ranking, since some agent in module $\mathcal{M}_{M-1}$ informs some agent in module $\mathcal{M}_M$, then some agent in module $\mathcal{M}_{M-1}$ must also inform agents in module $\mathcal{M}_{M-2}$, so there is bottom-up communication from module $\mathcal{M}_{M-1}$ to module $\mathcal{M}_{M-2}$. Further, module $\mathcal{M}_M$ dominates module $\mathcal{M}_{M-1}$ in the sender ranking and so since some agent in module $\mathcal{M}_{M-1}$ informs some agent in module $\mathcal{M}_{M-2}$, then some agent in module $\mathcal{M}_M$ must also inform some agent in module $\mathcal{M}_{M-2}$. So there is also bottom-up communication from module $\mathcal{M}_M$ to module $\mathcal{M}_{M-2}$.

Putting the results in steps 2 and 3 together, we have that for every pair of modules for which there is top-down communication, there are two unique pairs of modules for which there is bottom-up communication, so communication must be bottom up in aggregate. This proof has also established that if there exists top-down communication between any pair of modules in the optimal communication network, then aggregate communication is strictly bottom up (i.e., there are strictly more pairs of modules that engage in bottom-up communication than engage in top-down communication). Further, if there is no top-down communication between any pair of modules, then aggregate communication must also be strictly bottom-up, except for the case where there is some communication from an agent in each module to an agent in all other modules. ■

PROPOSITION 4. *If the optimal sender and receiver hierarchies are identical, the communication network has a core-periphery structure in which the core consists of the most cohesive modules.*

Proof of Proposition 4. Label the modules by their cohesion, with the most cohesive labeled $\mathcal{M}_1$ and the least cohesive labeled $\mathcal{M}_M$, that is $x_1 \geq x_2 \geq \cdots \geq x_M$. Find the highest $k \in \{2, \ldots, M\}$ such that some agent in module k informs some agent in module $\mathcal{M}_{k-1}$. We deal with the case where no such k exists below.

An agent either informs everyone in a module or no one in that module, and thus there exists some agent in module $\mathcal{M}_k$ who informs all agents in module $\mathcal{M}_{k-1}$ and, by the receiver ranking, all agents in modules $\mathcal{M}_1, \ldots, \mathcal{M}_{k-1}$. Since the sender ranking is the same as the receiver ranking, if an agent in module $\mathcal{M}_k$ informs all agents in modules $\mathcal{M}_1, \ldots, \mathcal{M}_{k-1}$, then all agents in any module $\mathcal{M}_m \in \{\mathcal{M}_1, \ldots, \mathcal{M}_{k-1}\}$ must inform all other agents in modules $\mathcal{M}_1, \ldots, \mathcal{M}_{k-1}$. Therefore there is full communication among agents in modules $\mathcal{M}_1, \ldots, \mathcal{M}_{k-1}$.

We next consider modules $\mathcal{M}_{k+1}, \ldots, \mathcal{M}_M$. By the above definition, there is no agent in module $\mathcal{M}_{k+1}$ who informs an agent in module $\mathcal{M}_k$. By the receiver ranking, there is no agent in module $\mathcal{M}_{k+1}$ who informs any agent in modules $\mathcal{M}_k, \mathcal{M}_{k+1}, \ldots, \mathcal{M}_M$ (aside from agents in his own module). Since the sender ranking is the same as the receiver ranking, if an agent in module $\mathcal{M}_{k+1}$ does not tell his state to agents in modules $\mathcal{M}_k, \mathcal{M}_{k+1}, \ldots, \mathcal{M}_M$, then any agent in module $\mathcal{M}_{k+2}, \ldots, \mathcal{M}_M$ does not inform any agents in modules $\mathcal{M}_k, \mathcal{M}_{k+1}, \ldots, \mathcal{M}_M$ (aside from agents in his own module). Therefore there is no communication across modules $\mathcal{M}_{k+1}, \ldots, \mathcal{M}_M$.

It remains to determine whether agents in module $\mathcal{M}_k$ are in the core or periphery. Suppose all agents in module $\mathcal{M}_{k-1}$ inform all agents in module $\mathcal{M}_k$. From the sender ranking, it then follows that all agents in modules $\mathcal{M}_1, \ldots, \mathcal{M}_{k-1}$ inform all agents in module $\mathcal{M}_k$. In this case, any agents in module $\mathcal{M}_k$ that inform agents in modules $\mathcal{M}_1, \ldots, \mathcal{M}_{k-1}$ are in the core. By the receiver ranking, agents in module $\mathcal{M}_k$ that do not inform agents in module $\mathcal{M}_{k-1}$ also do not inform any agents in modules $\mathcal{M}_{k+2}, \ldots, \mathcal{M}_M$, and so those agents are in the periphery.

Next, suppose some agents in module $\mathcal{M}_{k-1}$ do not inform agents in module $\mathcal{M}_k$. Then module $\mathcal{M}_k$ is part of the periphery. By the receiver ranking, since agents in module $\mathcal{M}_{k-1}$ do not inform agents in module $\mathcal{M}_k$, they also do not inform agents in modules with lower cohesion, $\mathcal{M}_{k+1}, \ldots, \mathcal{M}_M$. Then, by the sender ranking, agents in module $\mathcal{M}_k$ also do not communicate with agents in modules $\mathcal{M}_{k+1}, \mathcal{M}_{k+2}, \ldots, \mathcal{M}_M$ and so $\mathcal{M}_k$ is in the periphery.

If there is no $k \in \{2, \ldots, M\}$ such that some agent in module $\mathcal{M}_k$ informs an agent in module $\mathcal{M}_{k-1}$, then no agent in module $\mathcal{M}_2$ informs any agent in module $\mathcal{M}_1$, and by the sender and receiver rankings, no agent in modules $\mathcal{M}_2, \ldots, \mathcal{M}_M$ inform any agent outside their own module. If some agents in module $\mathcal{M}_1$ inform others, those agents in module $\mathcal{M}_1$ form the core, and all other agents inform no one outside their own module and form the periphery. If no agents inform anyone outside their own module, then all agents are in the periphery. ■

PROPOSITION 5. *Mirroring is optimal if and only if* $t \leq \min_{i\in\mathcal{N}} \bar{t}_i$, *where* $\bar{t}_i > 0$ *is the threshold degree of coupling above which it is optimal for agent* i *to tell his state to agents in modules other than his own and below which it is not. Adding modules to the production function decreases the threshold* $\bar{t}_i$, *as does increasing the module characteristics* $n_{m'}$ *or* $p_{m'}$ *for any* $m' \neq m(i)$.

Proof of Proposition 5. *Without loss of generality, we consider agent* 1 *and show there exists a threshold* $\bar{t}_1 > 0$ *such that profits are higher if agents in other modules are not informed about* θ_1 *if and only if* $t < \bar{t}_1$.

From Lemma 3, the expected revenue generated by agent 1 when some ℓ modules $\mathcal{M}_1, \ldots, \mathcal{M}_\ell$, know state θ_1, is

$$a_1^2\sigma_1^2 \left(\frac{1 - p_1 (n_1 - 2)}{(1 + p_1)(1 - p_1 (n_1 - 1))} + \frac{t^2 x_1^2 \sum_{j=2}^{\ell} n_j x_j}{(1 - t n_1 x_1)\left(1 - t\sum_{j=1}^{\ell} n_j x_j\right)} \right).$$

The first term, $a_1^2\sigma_1^2 \frac{1-p_1(n_1-2)}{(1+p_1)(1-p_1(n_1-1))}$, is the expected revenue when agent 1 and his own module $\mathcal{M}_1$ are informed about θ_1. The second term is the expected revenue of additionally informing modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$. Expected profits from informing modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$ are positive only if the expected per-node incremental revenue from informing those modules exceeds the cost,

$$\frac{a_1^2\sigma_1^2}{n_2 + \cdots + n_\ell} \frac{t^2 x_1^2 \sum_{j=1}^{\ell} n_j x_j}{(1 - t n_1 x_1)\left(1 - t\sum_{j=1}^{\ell} n_j x_j\right)} \geq \gamma.$$

The left-hand side of this inequality is zero if $t = 0$, and profit from informing this set of communities is negative. The terms in the numerator, $t^2 x_1^2 n_j x_j$, are increasing in t, and the terms in the denominator, $-tn_jx_j$, are decreasing in t, and so the left-hand side is increasing in t. Thus there exists some threshold that we denote by $\bar{t}_\ell$ such that when $t \leq \bar{t}_\ell$, profit is less than or equal to zero from informing this set of modules. Since we assume $t \leq p_m$ for all $m \in \{1, \ldots, M\}$, then $\bar{t}_\ell$ occurs either at the value of t for which the left-hand side of the inequality is equal to γ (and above such a threshold there is positive profit from informing that set of modules) or at $\min\{p_m : m \in \{1, \ldots, M\}\}$, whichever is the smaller of the two.

Let $\mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M)$ denote the power set of the modules $\{\mathcal{M}_2, \ldots, \mathcal{M}_M\}$. Let e denote an element of $\mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M)$, where e is a set of modules. Let $\bar{t}_e$ denote the threshold such that it is not profitable to inform the modules in e about θ_1 if and only if $t \leq \bar{t}_e$. Then it is not profitable to inform any other modules about θ_1 if and only if $t \leq \min\left\{\bar{t}_e : e \in \mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M) \backslash \emptyset\right\}$ and therefore $\bar{t}_1 = \min\left\{\bar{t}_e : e \in \mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M) \backslash \emptyset\right\}$.

We next show that $\bar{t}_1$ is decreasing in n_m and p_m. Terms with n_m enter the numerator on the left-hand side of the inequality above via the term $\frac{n_m x_m}{n_2 + \cdots + n_m + \cdots + n_\ell}$, which is increasing in n_m, and enter the denominator via the terms $-n_m x_m$, which are decreasing in n_m (since by assumption, $p_m \geq t$). Therefore the left-hand side of the inequality is increasing in n_m for all $m \neq 1$. It is also increasing in p_m. An increase in n_m or p_m lowers the threshold value of t at which the left-hand side of the inequality is equal to γ and therefore weakly lowers the threshold $\bar{t}_e$ for any $e \in \mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M)$. It follows that $\bar{t}_1 = \min\left\{\bar{t}_e : e \in \mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M) \backslash \emptyset\right\}$ cannot increase following an increase in n_m or p_m for all $m \neq 1$.

Next, consider an increase in the number of modules such that we add an additional module $\mathcal{M}_{M+1}$ without making changes to the existing modules $\{\mathcal{M}_2, \ldots, \mathcal{M}_M\}$. Then $\mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M) \subset \mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M, \mathcal{M}_{M+1})$ and hence

$$\min\left\{\bar{t}_e : e \in \mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M) \backslash \emptyset\right\} \subset \min\left\{\bar{t}_e : e \in \mathcal{P}(\mathcal{M}_2, \ldots, \mathcal{M}_M, \mathcal{M}_{M+1}) \backslash \emptyset\right\},$$

and therefore $\bar{t}_1$ cannot increase following an increase in the number of modules. ■

PROPOSITION 6. *When partial mirroring is optimal, the organization contains one cluster of modules, and the modules that form the cluster are the most cohesive ones.*

Proof of Proposition 6. We first show that when an organization partially mirrors the production function, the set of modules are partitioned such that only one subset contains more than one module. Suppose the set of modules are partitioned such that there are at least two subsets each containing at least two modules. We show a contradiction. Without loss of generality, we suppose one subset includes modules denoted $\mathcal{M}_2$ and $\mathcal{M}_3$ and possibly others, and another subset contains modules $\mathcal{M}_4$ and $\mathcal{M}_5$ and possibly others. By Proposition 2, since $\mathcal{M}_2$ informs $\mathcal{M}_3$ but not $\mathcal{M}_5$, it must be that $x_3 > x_5$, but since $\mathcal{M}_4$ informs $\mathcal{M}_5$ but not $\mathcal{M}_3$, then $x_5 > x_3$, which is a contradiction.

We next consider which modules are part of the subset that contains multiple modules when an organization partially mirrors the production function. Suppose the module labeled $\mathcal{M}_s$ is part of this subset. By Proposition 2, an agent i in module s tells his state to an agent j outside i's module if and only if $x_{m(j)} \geq \lambda_i$. The symmetric argument holds for all individuals within all modules in this subset. Therefore this subset contains the modules $\mathcal{M}_s$ with the highest values of x_s. ■

PROPOSITION 7. *If agents internalize only a fraction $\mu \in [0,1]$ of the needs to coordinate, an optimal communication network solves*

$$\max_{\boldsymbol{C}} \sum_{i=1}^{N} a_i \mathrm{Cov}\left(d_i^*, \theta_i\right) + (1-\mu) \sum_{i=1}^{N} \sum_{j=1}^{N} p_{ij} \mathrm{Cov}\left(d_i^*, d_j^*\right) - \gamma \sum_{i=1}^{N} \sum_{j=1}^{N} m_{ij} c_{ij},$$

where

$$\mathrm{Cov}\left(d_i^*, \theta_i\right) = a_i \sigma_i^2 \omega_{ii}\left(\boldsymbol{C}_i, \mu\right)$$

and

$$\mathrm{Cov}\left(d_i^*, d_j^*\right) = \sum_{s=1}^{N} a_s^2 \sigma_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right),$$

and where $\omega_{ij}\left(\boldsymbol{C}_j, \mu\right)$ denotes the ijth entry of $\left(\boldsymbol{I} - \left(\mathrm{diag}\, \boldsymbol{C}_j\right) \mu \boldsymbol{P} \left(\mathrm{diag}\, \boldsymbol{C}_j\right)\right)^{-1}$, and m_{ij} is a dummy variable that is equal to one if and only if agents i and j belong to different modules.

Proof of Proposition 7. Suppose agents internalize only fraction $\mu \in [0,1]$ of the need to coordinate. Then they act as if the need for coordination is μp_{ij} rather than p_{ij}. Their best response functions are therefore

$$d_i = a_i \theta_i + \sum_{j=1}^{N} \mu p_{ij} \mathrm{E}_i\left[d_j\right],$$

and so by Lemma 1, equilibrium decisions are given by

$$d_i^* = \sum_{j=1}^{N} a_j \omega_{ij}\left(\boldsymbol{C}_j, \mu\right) \theta_j \text{ for all } i \in \mathcal{N},$$

where $\omega_{ij}\left(\boldsymbol{C}_j, \mu\right)$ is the ijth entry of $\left(\boldsymbol{I} - \left(\mathrm{diag}\, \boldsymbol{C}_j\right) \mu \boldsymbol{P} \left(\mathrm{diag}\, \boldsymbol{C}_j\right)\right)^{-1}$.

We know from the argument preceding Lemma 2 that we can write equilibrium revenue as

$$\begin{aligned} r\left(d_1^*, \ldots, d_N^*\right) &= \sum_{i=1}^{N} a_i d_i^* \theta_i - \sum_{i=1}^{N} d_i^* \left(d_i^* - a_i \theta_i - \sum_{j=1}^{N} p_{ij} d_j^*\right) \\ &= \sum_{i=1}^{N} a_i d_i^* \theta_i - \sum_{i=1}^{N} \sum_{j=1}^{N} p_{ij} d_i^* \left(\mu \mathrm{E}_i\left[d_j^*\right] - d_j^*\right), \end{aligned}$$

where the second line follows from the first by replacing d_i^* inside the parentheses with its best response.

Taking expectations, we have

$$\mathrm{E}\left[r\left(d_1^*, \ldots, d_N^*\right)\right] = \sum_{i=1}^{N} a_i \mathrm{E}\left[d_i^* \theta_i\right] - \sum_{i=1}^{N} \sum_{j=1}^{N} p_{ij} \mathrm{E}\left[d_i^* \left(\mu \mathrm{E}_i\left[d_j^*\right] - d_j^*\right)\right].$$

Substituting equilibrium decisions into the second term, we have

$$\begin{aligned}
\mathrm{E}\left[d_i^* \left(\mu \mathrm{E}_i\left[d_j^*\right] - d_j^*\right)\right] &= \mathrm{E}\left[\sum_{s=1}^{N} \sum_{t=1}^{N} a_s a_t \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{jt}\left(\boldsymbol{C}_t, \mu\right) \left(\mu \mathrm{E}_i\left[\theta_t\right] - \theta_t\right) \theta_s\right] \\
&= \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \left(\mu \mathrm{E}_i\left[\theta_s\right] - \theta_s\right) \theta_s\right] \\
&= \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \left(\mathrm{E}_i\left[\theta_s\right] - \theta_s\right) \theta_s\right] \\
&\quad - (1-\mu) \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \mathrm{E}_i\left[\theta_s\right] \theta_s\right],
\end{aligned}$$

where the second equality holds by independence of the states. We can rewrite the second term in this last expression as

$$\begin{aligned}
&(1-\mu) \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \theta_s^2\right] \\
&+ (1-\mu) \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \left(\mathrm{E}_i\left[\theta_s\right] - \theta_s\right) \theta_s\right],
\end{aligned}$$

which allows us to write

$$\begin{aligned}
\mathrm{E}\left[d_i^* \left(\mu \mathrm{E}_i\left[d_j^*\right] - d_j^*\right)\right] &= \mu \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \left(\mathrm{E}_i\left[\theta_s\right] - \theta_s\right) \theta_s\right] \\
&\quad - (1-\mu) \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \theta_s^2\right].
\end{aligned}$$

The first term on the right-hand side is zero because of the binary communication structure: either i knows θ_s, in which case $\mathrm{E}_i\left[\theta_s\right] - \theta_s = 0$, or he does not, in which case $\omega_{is}\left(\boldsymbol{C}_s, \mu\right) = 0$.

The second term can be stated as

$$-(1-\mu) \mathrm{E}\left[\sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \theta_s^2\right] = -(1-\mu) \operatorname{Cov}\left(d_i^*, d_j^*\right).$$

To see why, note that for two equilibrium decisions

$$d_i^* = \sum_{s=1}^{N} a_s \omega_{is} \left(\boldsymbol{C}_s, \mu\right) \theta_s \text{ and } d_j^* = \sum_{t=1}^{N} a_t \omega_{jt} \left(\boldsymbol{C}_s, \mu\right) \theta_t,$$

we have $\mathrm{E}\left[d_i^*\right] = \mathrm{E}[d_j^*] = 0$, and so covariance is given by

$$\begin{aligned} \mathrm{Cov}\left(d_i^*, d_j^*\right) &= \mathrm{E}\left[d_i^* d_j^*\right] = \mathrm{E}\left[\sum_{s=1}^{N}\sum_{t=1}^{N} a_s a_t \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{jt}\left(\boldsymbol{C}_t, \mu\right) \theta_s \theta_t\right] \\ &= \sum_{s=1}^{N} a_s^2 \omega_{is}\left(\boldsymbol{C}_s, \mu\right) \omega_{js}\left(\boldsymbol{C}_s, \mu\right) \sigma_s^2. \end{aligned}$$

We therefore have that

$$\mathrm{E}\left[d_i^* \left(\mu \mathrm{E}_i\left[d_j^*\right] - d_j^*\right)\right] = -\left(1-\mu\right) \mathrm{Cov}\left(d_i^*, d_j^*\right).$$

Substituting this expression back, we have

$$\begin{aligned} \mathrm{E}\left[r\left(d_1^*, \ldots, d_N^*\right)\right] &= \sum_{i=1}^{N} a_i \mathrm{E}\left[d_i^* \theta_i\right] + (1-\mu) \sum_{i=1}^{N}\sum_{j=1}^{N} p_{ij} \mathrm{Cov}\left(d_i^*, d_j^*\right) \\ &= \sum_{i=1}^{N} \left(a_i \mathrm{Cov}\left(d_i^*, \theta_i\right) + (1-\mu) \sum_{j=1}^{N} p_{ij} \mathrm{Cov}\left(d_i^*, d_j^*\right) \right), \end{aligned}$$

which establishes the result. ■

PROPOSITION 8. *As long as the production network* $\boldsymbol{P}$ *satisfies* $p_{ii} = 0$, $p_{ij} = p_{ji}$, *and* $\sum_{j=1}^{N} p_{ij} < 1$, *optimal communication networks* $\boldsymbol{C}^*$ *are increasing in the value of autonomous adaptation* $a_i^2 \sigma_i^2$ *and the needs for coordination* p_{ij} *for all* $i, j \in \mathcal{N}$, *and decreasing in communication costs* γ.

Proof of Proposition 8. For general production networks $\boldsymbol{P}$ satisfying $p_{ii} = 0$, $p_{ij} = p_{ji}$, and $\sum_{j=1}^{N} p_{ij} < 1$, Lemmas 1 and 2 and Proposition 1 continue to hold. As long as $\omega_{ii}\left(\boldsymbol{C}_i\right)$ is supermodular in $\boldsymbol{C}_i$, then the principal's objective for the subproblem involving who should agent i inform about θ_i is supermodular in $\boldsymbol{C}_i$ and exhibits increasing differences in $\left(a_i^2 \sigma_i^2, \left\{p_{ij}\right\}_{ij}, \left\{c_{ij}\right\}_{ij}, -\gamma\right)$, so the comparative statics results follow from Topkis's theorem. It remains, therefore, to show that $\omega_{ii}\left(\boldsymbol{C}_i\right)$ is supermodular in $\boldsymbol{C}_i$.

To show that $\omega_{ii}\left(\cdot\right)$ is supermodular, let $\mathcal{J} \subset \mathcal{N}$ denote a subset of agents, and denote by $\mathbf{c}\left(\mathcal{J}\right)$ the $1 \times N$ vector with jth element equal to one if $j \in \mathcal{J}$ and equal to zero otherwise. We will show that the incremental value of informing agent 1 about θ_i is higher when agent 2 knows θ_i than when she does not. Take $\mathcal{J}$ to be a set of nodes that are informed throughout the exercise.

Denote by $\boldsymbol{P}(\mathcal{J}) = (\operatorname{diag}\mathbf{c}(\mathcal{J}))\,\boldsymbol{P}\,(\operatorname{diag}\mathbf{c}(\mathcal{J}))$ the subset of the production network consisting of the nodes j for which the jth element of $\mathbf{c}(\mathcal{J})$ is equal to one. Then

$$\Delta^k \equiv \boldsymbol{P}(\mathcal{J}\cup\{1,2\})^k - \boldsymbol{P}(\mathcal{J}\cup\{2\})^k - \left(\boldsymbol{P}(\mathcal{J}\cup\{1\})^k - \boldsymbol{P}(\mathcal{J})^k\right)$$

is the matrix whose ijth element is the value of the additional walks of length k from informing agent 1 when agents $\mathcal{J}\cup\{2\}$ are informed relative to when only agents $\mathcal{J}$ are informed. Since informing agent 1 adds more walks of all lengths to $\boldsymbol{P}(\mathcal{J}\cup\{2\})$ than it does to $\boldsymbol{P}(\mathcal{J})$, it follows that every element of Δ^k is nonnegative. Since this argument holds for all k, we have that the iith element of

$$\sum_{k=1}^{\infty}\Delta^k = \sum_{k=1}^{\infty}\boldsymbol{P}(\mathcal{J}\cup\{1,2\})^k - \sum_{k=1}^{\infty}\boldsymbol{P}(\mathcal{J}\cup\{2\})^k - \left(\sum_{k=1}^{\infty}\boldsymbol{P}(\mathcal{J}\cup\{1\})^k - \sum_{k=1}^{\infty}\boldsymbol{P}(\mathcal{J})^k\right)$$

is nonnegative. Recall that $\omega_{ii}(\mathbf{c}(\mathcal{J}))$ is the iith element of $(\boldsymbol{I}-\boldsymbol{P}(\mathcal{J}))^{-1} = \boldsymbol{I} + \sum_{k=1}^{\infty}\boldsymbol{P}(\mathcal{J})^k$. We therefore have that

$$\omega_{ii}(\mathbf{c}(\mathcal{J}\cup\{1,2\})) - \omega_{ii}(\mathbf{c}(\mathcal{J}\cup\{2\})) \geq \omega_{ii}(\mathbf{c}(\mathcal{J}\cup\{1\})) - \omega_{ii}(\mathbf{c}(\mathcal{J})),$$

so $\omega_{ii}(\cdot)$ has increasing differences in c_{i1} and c_{i2}. The choice of agents 1 and 2 was immaterial in this argument, and so $\omega_{ii}(\cdot)$ has increasing differences in c_{ij} and c_{ik} for all $j,k\neq i$ and is therefore supermodular. ■

PROPOSITION 9. *If the production network is symmetric and has a nested, modular structure, optimal communication networks are characterized by a threshold $\kappa\in\{1,2,\ldots,k\}$. Agent i tells his state to agent j if and only if they belong to the same level-κ module.*

Proof of Proposition 9. Since all the agents are symmetric, let us consider agent 1 and ask who agent 1 should inform about θ_1. By supermodularity and symmetry of the levels, if agent 1 informs any agent in his level-s module who is not in his level-$s-1$ module, for some $s\leq k$, then he does even better by informing all nodes in his level-s module who are not in his level-$s-1$ module. Therefore, it is without loss of optimality to consider communication networks in which agent 1 either informs everyone in his level-s module, or he informs everyone in his level-s module except those in his level-$s-r$ module for some $1\leq r\leq s-1$. We proceed in two steps.

Step 1: *Show that agent 1's per-node incremental revenues are higher if he informs everyone in his level-s module than if he informs everyone in his level-s module except those in his level-$s-1$ module, for any $s\leq k$.*

For this step, we will consider two communication networks: (i.) Agent 1 informs everyone in the other $n-1$ level-$s-1$ modules but no one in his own level-$s-1$ module; (ii.) Agent 1 informs everyone in his own level-$s-1$ module and everyone in $n-2$ other level-$s-1$ modules and one agent in the remaining level-$s-1$ module. We will show that agent 1's revenue is higher in (*ii.*). It follows by supermodularity and symmetry of the levels that his revenues are higher still if he informs *everyone* in his level-s module.

To make this argument, we proceed as follows. Suppose agent 1 informs the agents specified in (*ii.*): everyone in his own level-$s-1$ module and everyone in $n-2$ other level-$s-1$ modules and one agent in the remaining level-$s-1$ module. Label the other agents in agent 1's level-$s-1$ module and everyone in the $n-2$ other level-$s-1$ modules consecutively by $\left\{2,3,\ldots,(n-1)\,n^{s-1}\right\}$. Label the single agent in the remaining level-$s-1$ module as agent z. Notice that the revenue generated by agent z in communication network (*ii.*) is the same as the revenue generated by agent 1 in communication network (*i.*), so it will suffice to show that the value of all walks from agent 1 back to agent 1 exceed the value of all walks from agent z back to agent z. We will show this is the case by matching walks from agent 1 back to agent 1 with walks from agent z back to agent z.

We first introduce some notation that will help with this walk-matching exercise. Let $\mathcal{W}_\tau(i,j)$ be the set of walks of length τ from i to j, and denote a generic element of this set by $i.\overset{\tau}{.}.j$. Denote by $ij,j\overset{\tau-2}{\cdots}\ell,\ell m$ a walk of length τ starting with link ij, followed by a walk of length $\tau-2$ that starts at node j and ends at node ℓ, and ending with link ℓm. Let $\mathcal{W}(i,j)=\cup_{\tau=1}^{\infty}\mathcal{W}_\tau(i,j)$ be the set of walks from i to j with generic element $i\ldots j$, and denote the value of $i\ldots j$ by $\hat{v}(i\ldots j)$. It will be useful to keep in mind that the function $\hat{v}(\cdot)$ has the property that $\hat{v}(i\ldots j,j\ldots\ell)=\hat{v}(i\ldots j)\,\hat{v}(j\ldots\ell)$.

The goal is to find a bijection $H:\mathcal{W}(z,z)\to\mathcal{W}(1,1)$ such that $\hat{v}(H(z\ldots z))\geq\hat{v}(z\ldots z)$ for all $z\ldots z\in\mathcal{W}(z,z)$. To this end, let us partition the sets $\mathcal{W}_\tau(z,z)$ and $\mathcal{W}_\tau(1,1)$ as follows:

$$\mathcal{W}_\tau(z,z)=\left\{zi,i\overset{\tau-2}{\cdots}j,jz\,\middle|\,i,j\neq1\right\}\cup\left\{zi,i\overset{\tau-2}{\cdots}1,1z\,\middle|\,i\neq1\right\}\cup\left\{z1,1\overset{\tau-2}{\cdots}1,1z\right\}\cup\left\{z1,1\overset{\tau-2}{\cdots}j,jz\,\middle|\,j\neq1\right\},$$

and

$$\mathcal{W}_\tau(1,1)=\left\{1i,i\overset{\tau-2}{\cdots}j,j1\,\middle|\,i,j\neq z\right\}\cup\left\{1z,zi,i\overset{\tau-2}{\cdots}1\,\middle|\,i\neq1\right\}\cup\left\{1z,z1,1\overset{\tau-2}{\cdots}1\right\}\cup\left\{1j,j\overset{\tau-2}{\cdots}z,z1\,\middle|\,j\neq z\right\}.$$

First, let $H\left(zi,i\overset{\tau-2}{\cdots}j,jz\right)=1i,i\overset{\tau-2}{\cdots}j,j1$ for all $i,j\neq1,z$ and $i\overset{\tau-2}{\cdots}j\in\mathcal{W}_{\tau-2}(i,j)$. That is, replace the initial link zi with $1i$ and the terminal link jz with $j1$, and hold the walk $i\overset{\tau-2}{\cdots}j$ of length $\tau-2$ constant. Notice that $\hat{v}\left(H\left(zi,i\overset{\tau-2}{\cdots}j,jz\right)\right)\geq\hat{v}\left(zi,i\overset{\tau-2}{\cdots}j,jz\right)$ since the value of the links zi and jz are t_s, and the value of the links $1i$ and $j1$ are greater than t_s. Second, let $H\left(zi,i\overset{\tau-2}{\cdots}1,1z\right)=1z,zi,i\overset{\tau-2}{\cdots}1$ for all $i\neq1$, and $i\overset{\tau-2}{\cdots}1\in\mathcal{W}_{\tau-2}(i,1)$. That is, for any walk from z back to z that ends with the link $1z$, start that walk instead with the link $1z$ and continue until it returns to node 1. Notice that $\hat{v}\left(H\left(zi,i\overset{\tau-2}{\cdots}1,1z\right)\right)=\hat{v}\left(zi,i\overset{\tau-2}{\cdots}1,1z\right)$ since the walks contain the same set of links. Third, let $H\left(z1, 1\overset{\tau-2}{\cdots}1, 1z\right) = 1z, z1, 1\overset{\tau-2}{\cdots}1$ for all $1\overset{\tau-2}{\cdots}1 \in \mathcal{W}_{\tau-2}(1,1)$. Notice that $\hat{v}\left(H\left(z1, 1\overset{\tau-2}{\cdots}1, 1z\right)\right) = \hat{v}\left(z1, 1\overset{\tau-2}{\cdots}1, 1z\right)$ since the walks contain the same links.

It remains to match walks of the form $z1, 1\overset{\tau-2}{\cdots}j, jz$, $j \neq 1$ to walks of the form $1j, j\overset{\tau-2}{\cdots}z, z1$, $j \neq z$. Note that $z1, 1\overset{\tau-2}{\cdots}j, jz$ can be written as $z\overset{\tau_1}{\cdots}\ell, \ell i, i\overset{\tau_2}{\cdots}j, jz$, where $z\overset{\tau_1}{\cdots}\ell$ is the longest initial subwalk that alternates only between z and 1 (e.g., for the walk $z1, 1z, z1, 1z, z2, 23, 3z$, we would have $\tau_1 = 4$ and $\tau_2 = 1$). There are two cases: $\ell = 1$ and $\ell = z$. For $\ell = z$, let $H\left(z\overset{\tau_1}{\cdots}z, zi, i\overset{\tau_2}{\cdots}j, jz\right) = 1i, i\overset{\tau_2}{\cdots}j, j1, 1\overset{\tau_1}{\cdots}1$, where $1\overset{\tau_1}{\cdots}1$ alternates between 1 and z. In this case,

$$\frac{\hat{v}\left(H\left(z\overset{\tau_1}{\cdots}z, zi, i\overset{\tau_2}{\cdots}j, jz\right)\right)}{\hat{v}\left(z\overset{\tau_1}{\cdots}z, zi, i\overset{\tau_2}{\cdots}j, jz\right)} = \frac{\hat{v}(1i)\,\hat{v}\left(i\overset{\tau_2}{\cdots}j\right)\hat{v}(j1)\,\hat{v}\left(1\overset{\tau_1}{\cdots}1\right)}{\hat{v}\left(z\overset{\tau_1}{\cdots}z\right)\hat{v}(zi)\,\hat{v}\left(i\overset{\tau_2}{\cdots}j\right)\hat{v}(jz)} = \frac{\hat{v}(1i)\,\hat{v}(j1)}{\hat{v}(zi)\,\hat{v}(jz)} \geq 1,$$

where the first equality uses the property that $\hat{v}(i\ldots, j, j\ldots, \ell) = \hat{v}(i\ldots j)\,\hat{v}(j\ldots\ell)$, and the second equality uses the fact that $\hat{v}\left(1\overset{\tau_1}{\cdots}1\right) = \hat{v}\left(z\overset{\tau_1}{\cdots}z\right)$ because both are equal-length walks between z and 1. The inequality follows because $\hat{v}(1i) \geq \hat{v}(zi)$ for any i. For $\ell = 1$, let $H\left(z\overset{\tau_1}{\cdots}1, 1i, i\overset{\tau_2}{\cdots}j, jz\right) = 1i, i\overset{\tau_2}{\cdots}j, jz, z\overset{\tau_1}{\cdots}1$. In this case, $\hat{v}\left(H\left(z\overset{\tau_1}{\cdots}1, 1i, i\overset{\tau_2}{\cdots}j, jz\right)\right) = \hat{v}\left(z\overset{\tau_1}{\cdots}1, 1i, i\overset{\tau_2}{\cdots}j, jz\right)$.

The function H is a bijection $H : \mathcal{W}(z,z) \to \mathcal{W}(1,1)$ such that $\hat{v}(H(z\ldots z)) \geq \hat{v}(z\ldots z)$, and we therefore have that agent 1's revenues are higher if he informs everyone in his level-s module than if he informs everyone in his level-s module except those in his level-$s-1$ module.

Step 2: *Show that agent 1's per-node incremental revenues are higher if he informs everyone in his level-$s+r$ module than if he informs everyone in his level-$s+r$ module except those in his level-$s-1$ module, for $r \geq 1$.*

For this step, we will consider two communication networks: (i.) Agent 1 informs everyone in his level-$s+r$ module, for $r \geq 1$, but no one in his own level-$s-1$ module; (ii.) Agent 1 informs everyone in his level-$s+r$ module, except for one level-$s-1$ module, not his own, within his level-s module in which he informs only one agent, which we will label as agent z. As in step 1, the revenue generated by agent z in communication network (ii.) is the same as the revenue generated by agent 1 in communication network (i.).

We can again apply the same bijection H constructed in step 1, which establishes that the value of all walks from agent 1 back to agent 1 exceed the value of all walks from agent z back to agent z. By supermodularity and level symmetry, agent 1's per-node incremental revenues are higher still if he informs all agents in his level-$s-1$ module. This implies that agent 1's per-node incremental revenues are higher if he informs everyone in his level-$s+r$ module than if he informs everyone in his level-$s+r$ module except those in his level-$s-1$ module, for $r \geq 1$. Since s and r are arbitrary, it follows that there is some κ such that in any optimal communication network, agent 1 informs agent j if and only if they belong to the same level-κ module. ■

Appendix B: Equal Communication Costs

This appendix examines the case in which each communication link costs γ, whether that communication link is within or across modules. Proposition B1 shows that if agents do not inform their own module, optimal communication networks still follow a threshold communication rule analogous to that described in Proposition 2 for the main model. To this end, denote by $\tilde{R}_1(\boldsymbol{C}_1(\ell))$ agent 1's expected revenue if he tells his state to agents in modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$ but not to agents in his own module $\mathcal{M}_1$. The following lemma derives an expression for $\tilde{R}_1(\boldsymbol{C}_1(\ell))$ in terms of the model primitives.

LEMMA B1. *Suppose agent 1 tells his state to all agents in modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$ for $\ell \in \{2, \ldots, M\}$ but not to agents in his own module $\mathcal{M}_1$. Agent 1's expected revenue is then given by*

$$\tilde{R}_1(\boldsymbol{C}_1(\ell)) = a_1^2 \sigma_1^2 \left(1 + \frac{t^2 \left(\sum_{m=1}^{\ell} n_m x_m - n_1 x_1 \right)}{1 - (t + t^2) \sum_{m=2}^{\ell} n_m x_m} \right),$$

where

$$x_m = \frac{1}{1 - (n_m - 1) p_m + n_m t} \text{ for } m = 1, \ldots, M.$$

Proof of Lemma B1. The proof of this lemma parallels the proof of Lemma 3, but for the case where agents in module $\mathcal{M}_1$ do not know θ_1. The value $\tilde{\omega}_{11}(\boldsymbol{C}_1(\ell))$ is the sum of the values of all walks from node 1 back to itself on the subgraph of the production network consisting of nodes in modules whose agents know state θ_1. We first derive a recursive representation of $\tilde{\omega}_{11}(\boldsymbol{C}_1(\ell))$. Denote this value by $\tilde{v}$. Next, let $\tilde{v}_k$ be the sum of the values of all walks from a node in module k to node 1 on this same subgraph. These values can be written as a system of equations

$$\begin{aligned}
\tilde{v} &= 1 + t n_2 \tilde{v}_2 + \cdots + t n_\ell \tilde{v}_\ell \\
\tilde{v}_2 &= t \tilde{v} + p_2 (n_2 - 1) \tilde{v}_2 + t n_3 \tilde{v}_3 + \cdots + t n_\ell \tilde{v}_\ell \\
&\vdots \\
\tilde{v}_\ell &= t \tilde{v} + t n_2 \tilde{v}_2 + t n_3 \tilde{v}_3 + \cdots + p_\ell (n_\ell - 1) \tilde{v}_\ell.
\end{aligned}$$

Compared to the system of equations derived in the first step of Lemma 3, here, the terms with

v_1 are eliminated. We can write this system of ℓ equations in matrix form

$$\begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & -tn_2 & \cdots & tn_\ell \\ -t & 1 - p_2(n_2 - 1) & \cdots & tn_\ell \\ \vdots & \vdots & \ddots & \vdots \\ -t & -tn_2 & \cdots & 1 - p_\ell(n_\ell - 1) \end{bmatrix} \begin{bmatrix} \tilde{v} \\ \tilde{v}_2 \\ \vdots \\ \tilde{v}_\ell \end{bmatrix}.$$

We denote this $\ell \times \ell$ matrix by $\tilde{\boldsymbol{Q}}(\ell)$. Then $\tilde{v}$ is the $(1,1)$ element of the inverse matrix $\tilde{\boldsymbol{Q}}(\ell)^{-1}$, and by the definition of a matrix inverse, $\tilde{v} = \det \tilde{\boldsymbol{V}}(\ell) / \det \tilde{\boldsymbol{Q}}(\ell)$, where $\tilde{\boldsymbol{V}}(\ell)$ is the matrix obtained by removing the first row and column of $\tilde{\boldsymbol{Q}}(\ell)$.

Next, we derive $\det \tilde{\boldsymbol{Q}}(\ell)$ using the result for the determinant of a block matrix that $\det \begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{C} & \boldsymbol{D} \end{bmatrix} = \det(\boldsymbol{A}) \det\left(\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}\right)$. Partition the matrix $\tilde{\boldsymbol{Q}}(\ell)$ such that $\boldsymbol{A} = 1$, and $\boldsymbol{B}$, $\boldsymbol{C}$, and $\boldsymbol{D}$ are defined accordingly. Then

$$\boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B} = t^2 \begin{bmatrix} n_2 & \cdots & n_\ell \\ \vdots & \ddots & \vdots \\ n_2 & \cdots & n_\ell \end{bmatrix}.$$

Define $x_s^{-1} = 1 - p_s(n_s - 1) + tn_s$ and $z = t(1+t)$. Then

$$\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B} = \begin{bmatrix} x_2^{-1} - zn_2 & \cdots & -zn_\ell \\ \vdots & \ddots & \vdots \\ -zn_2 & \cdots & x_\ell^{-1} - zn_\ell \end{bmatrix}.$$

We can immediately use our derivation from Lemma 3, step 2 to find $\det\left(\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}\right) = \frac{1}{x_2 \cdots x_\ell}\left(1 - z\sum_{j=2}^{\ell} n_j x_j\right)$ for all $\ell \geq 2$. Since $\det \boldsymbol{A} = 1$, we have

$$\det \tilde{\boldsymbol{Q}}(\ell) = \frac{1}{x_2 \cdots x_\ell}\left(1 - t(1+t)\sum_{j=2}^{\ell} n_j x_j\right) \text{ for all } \ell \geq 2.$$

Similarly, we can use the derivation from Lemma 3, step 3 to show that

$$\det \tilde{\boldsymbol{V}}(\ell) = \frac{1}{x_2 \cdots x_\ell}\left(1 - t\sum_{j=2}^{\ell} n_j x_j\right) \text{ for all } \ell \geq 2.$$

The value $\tilde{\omega}_{11}(\mathbf{C}_1(\ell))$ is therefore

$$\tilde{\omega}_{11}(\boldsymbol{C}_1(\ell)) = \frac{\det \tilde{\boldsymbol{V}}(\ell)}{\det \tilde{\boldsymbol{Q}}(\ell)} = 1 + \frac{t^2 \sum_{j=2}^{\ell} n_j x_j}{1 - (t + t^2)\sum_{j=2}^{\ell} n_j x_j},$$

and the expression in the statement of the lemma follows immediately. ■

PROPOSITION B1. *When agents do not inform their own module, optimal communication is characterized by N thresholds $\overline{\lambda}_i \geq 0$, one for each agent $i \in \mathcal{N}$. Agent i tells his state to agent j, who is not in his module, if and only if $x_{m(j)} \geq \overline{\lambda}_i$. The threshold $\overline{\lambda}_i$ is decreasing in $a_i^2\sigma_i^2$ and increasing in communication costs γ. Also, $\overline{\lambda}_i \geq \lambda_i$, where λ_i is the threshold for agent i described in Proposition 2.*

Proof of Proposition B1. The proof of this lemma parallels the proof of Proposition 2, but for the case where agents in module $\mathcal{M}_1$ do not know θ_1. Denote $n_1 = 1$ and $x_1 = 1/(1+t)$. Suppose agents in modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$ are informed of θ_1, for some $\ell \geq 2$. Using the result from Lemma B1, the expected per-node incremental revenue of informing one more module, $\mathcal{M}_{\ell+1}$, about state θ_1, given that modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$ are informed, is

$$a_1^2\sigma_1^2 \frac{t^2 x_1^2 x_{\ell+1}}{\left(1 - t\sum_{j=1}^{\ell} n_j x_j\right)\left(1 - t\sum_{j=1}^{\ell+1} n_j x_j\right)}.$$

Analogously, the expected per-node incremental revenue of informing modules $\mathcal{M}_{\ell+1}$ and $\mathcal{M}_{\ell+2}$ about state θ_1, given that modules $\mathcal{M}_2, \ldots, \mathcal{M}_\ell$ are informed, is

$$\frac{1}{n_{\ell+1}+n_{\ell+2}} a_1^2\sigma_1^2 \frac{t^2 x_1^2 \left(n_{\ell+1}x_{\ell+1} + n_{\ell+2}x_{\ell+2}\right)}{\left(1 - t\sum_{j=1}^{\ell} n_j x_j\right)\left(1 - t\sum_{j=1}^{\ell+2} n_j x_j\right)}.$$

Note that this is the same expression for the case where module $\mathcal{M}_1$ is also informed, but in this case, we are setting $n_1 = 1$ and $x_1 = 1/(1+t)$.

Then the expected per-node incremental revenue of additionally informing modules $\mathcal{M}_{\ell+1}$ and $\mathcal{M}_{\ell+2}$ minus the expected per-node incremental revenue of additionally informing only module $\mathcal{M}_{\ell+1}$ about θ_1 is also the same as before

$$\frac{a_1^2\sigma_1^2 x_1^2 t^2 n_{\ell+2}\left[\left(x_{\ell+2} - x_{\ell+1}\right)\left(1 - t\sum_{j=1}^{\ell+1} n_j x_j\right) + t\left(n_{\ell+1}+n_{\ell+2}\right)x_{\ell+1}x_{\ell+2}\right]}{\left(n_{\ell+1}+n_{\ell+2}\right)\left(1 - t\sum_{j=1}^{\ell} n_j x_j\right)\left(1 - t\sum_{j=1}^{\ell+1} n_j x_j\right)\left(1 - t\sum_{j=1}^{\ell+2} n_j x_j\right)},$$

but with $n_1 = 1$ and $x_1 = 1/(1+t)$. Therefore the same argument for a threshold rule applies, and modules are ranked in the same way using the same statistic x_s. An analogous argument shows the threshold $\overline{\lambda}_i$ is decreasing in $a_i^2\sigma_i^2$ and increasing in γ. The result that $\overline{\lambda}_i \geq \lambda_i$ follows from the above result and Proposition 2, which states that λ_i is decreasing in n_i. ■

Appendix C: Imperfect Communication

This appendix derives the principal's problem when communication is imperfect. Given a communication network $\boldsymbol{C}$, suppose that when agent i communicates his state θ_i to all agents for which $c_{ij} = 1$, his communication is effective with probability q and ineffective with probability $1 - q$. If his communication is effective, then all agents j with $c_{ij} = 1$ learn θ_i, and if his communication is ineffective, then all agents j with $c_{ij} = 1$ receive an uninformative null signal. Suppose the realization of whether communication from agent i is effective is common knowledge. As in the main model, communication links to other agents in one's module are costless.

The timing is: First, the principal designs the communication network $\boldsymbol{C}$. Then the agents learn their states and communicate them to the other agents as specified by the network. Next, agents observe whose communication is effective, and they learn the states that were communicated successfully to them. Finally, agents simultaneously make their decisions, payoffs are realized, and the game ends.

Given a communication network $\mathbf{C}$ and the realization of whose communication was effective, denote by $\tilde{\boldsymbol{C}}$ the network that describes who is informed of which state. That is, $\tilde{c}_{ij} = 1$ if $c_{ij} = 1$ and agent i's communication was effective. Then the principal's problem is to design the optimal communication network that solves

$$\max_{\boldsymbol{C}} \mathrm{E}\left[r\left(d_1^*\left(\tilde{\boldsymbol{C}}\right), \ldots, d_N^*\left(\tilde{\boldsymbol{C}}\right)\right) \middle| \boldsymbol{C} \right] - \gamma \sum_{i=1}^{N} \sum_{j=1}^{N} m_{ij} c_{ij} \text{ subject to } c_{ii} = 1 \text{ for all } i \in \mathcal{N},$$

where $d_i^*\left(\tilde{\boldsymbol{C}}\right)$ denotes agent i's equilibrium decision given the realization of $\tilde{\boldsymbol{C}}$ and $\theta_1, \ldots, \theta_N$, and the expectation is taken over the realizations of $\tilde{\boldsymbol{C}}$ and $\theta_1, \ldots, \theta_N$. The next proposition shows that the principal's problem when communication is imperfect is equivalent to the principal's problem in the main model, except that the cost of each communication link is scaled up by a factor of $1/q$.

PROPOSITION C1. *An optimal communication network solves the principal's problem if and only if it solves the N independent subproblems*

$$\max_{\boldsymbol{C}_i} R_i\left(\boldsymbol{C}_i\right) - \tilde{\gamma} \sum_{j=1}^{N} m_{ij} c_{ij} \text{ subject to } c_{ii} = 1 \text{ for all } i \in \mathcal{N},$$

where m_{ij} is a dummy variable that is equal to one if and only if agents i and $j \neq i$ belong to different modules, and $\tilde{\gamma} = \gamma/q$.

Proof of Proposition C1. To establish this proposition, which parallels Proposition 1, we have to argue that a version of Lemmas 1 and 2 hold. First, note that the proof of Lemma 1 depended only on which agent knew which state. This implies that given a matrix $\tilde{\boldsymbol{C}}$ of who is informed of which state, equilibrium decisions are unique and given by

$$d_i^*\left(\tilde{\boldsymbol{C}}\right) = \sum_{j=1}^{N} a_j \omega_{ij}\left(\tilde{\boldsymbol{C}}_j\right)\theta_j \text{ for all } i \in \mathcal{N},$$

where $\omega_{ij}\left(\tilde{\boldsymbol{C}}_j\right)$ denotes the ijth entry of $\left(\boldsymbol{I} - \left(\operatorname{diag}\tilde{\boldsymbol{C}}_j\right)\boldsymbol{P}\left(\operatorname{diag}\tilde{\boldsymbol{C}}_j\right)\right)^{-1}$.

To see why a version of Lemma 2 holds, note that

$$\begin{aligned}
\mathrm{E}^{\tilde{\boldsymbol{C}},\theta}\left[r\left(d_1^*\left(\tilde{\boldsymbol{C}}\right),\ldots,d_N^*\left(\tilde{\boldsymbol{C}}\right)\right)\middle|\,\boldsymbol{C}\right] &= \mathrm{E}^{\tilde{\boldsymbol{C}}}\left[\mathrm{E}^{\theta}\left[r\left(d_1^*\left(\tilde{\boldsymbol{C}}\right),\ldots,d_N^*\left(\tilde{\boldsymbol{C}}\right)\right)\middle|\,\tilde{\boldsymbol{C}}\right]\middle|\,\boldsymbol{C}\right] \\
&= \mathrm{E}^{\tilde{\boldsymbol{C}}}\left[\sum_{i=1}^{N} a_i^2\sigma_i^2\omega_{ii}\left(\tilde{\boldsymbol{C}}_i\right)\middle|\,\boldsymbol{C}\right] \\
&= \sum_{i=1}^{N} a_i^2\sigma_i^2\mathrm{E}^{\tilde{\boldsymbol{C}}}\left[\omega_{ii}\left(\tilde{\boldsymbol{C}}_i\right)\middle|\,\boldsymbol{C}\right],
\end{aligned}$$

where the superscript of the expectation denotes which variable is being integrated over. The first equality holds by the law of iterated expectations, and the second equality holds by Lemma 2, which applies realization-by-realization of $\tilde{\boldsymbol{C}}$.

Finally, note that $\mathrm{E}^{\tilde{\boldsymbol{C}}}\left[\omega_{ii}\left(\tilde{\boldsymbol{C}}_i\right)\middle|\,\boldsymbol{C}\right] = (1-q) + q\omega_{ii}(\boldsymbol{C}_i)$. With probability $1-q$, no one other than agent i is informed about θ_i, so $\omega_{ii}\left(\tilde{\boldsymbol{C}}_i\right)$ is equal to one. With probability q, all agents j with $c_{ij} = 1$ are informed about θ_i, so $\omega_{ii}\left(\tilde{\boldsymbol{C}}_i\right) = \omega_{ii}(\boldsymbol{C}_i)$. The principal's objective is therefore separable across agents i, and her objective for agent i is to

$$\max_{\boldsymbol{C}_i} a_i^2\sigma_i^2(1-q) + a_i^2\sigma_i^2 q\omega_{ii}(\boldsymbol{C}_i) - \gamma\sum_{j=1}^{N} m_{ij}c_{ij}.$$

The first term is independent of $\boldsymbol{C}_i$, and the second term is just $qR_i(\boldsymbol{C}_i)$, so solving this problem is equivalent to solving

$$\max_{\boldsymbol{C}_i} R_i(\boldsymbol{C}_i) - \frac{\gamma}{q}\sum_{j=1}^{N} m_{ij}c_{ij},$$

which establishes the result. ■