Organizing Modular Production*

Niko Matouschek
Northwestern University

Michael Powell
Northwestern University

Bryony Reich
Northwestern University

June 23, 2023

Abstract

Products are increasingly made by assembling separately produced modules. Motivated by the notion that a firm's production function drives its organization, we explore how modular production shapes a firm's communication structure. Decisions are partitioned into modules and require closer coordination within modules than across. Each agent knows the state his decision must be adapted to. The principal decides whom each agent tells about his state, given that each communication link comes at a cost. We show that optimal communication networks follow a simple threshold rule and exhibit the *threshold property*. Comparative statics corroborate the *Mirroring Hypothesis*.

Keywords: organization, communication, network, modular

JEL classifications: D23, D85, L23

*We thank Heski Bar-Isaac, Hector Chade, Krishna Dasaratha, Wouter Dessein, Matt Elliott, Itay Fainmesser, Ben Golub, Matt Jackson, Willemien Kets, Elliot Lipnowski, 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 workshop 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

Modular production used to be the exception, now it is the norm. Ever since IBM introduced the first modular computer in 1964, firms in a wide range of industries have designed modular versions of their products (Baldwin and Clark 2000). Boeing's Dreamliner is one prominent example.[1] Other examples range from smartphones to residential homes to software programs.[2] Nowadays so many products are made by assembling separately produced modules that our times have been called the *Modular Age* (Garud, Kumaraswamy, and Langlois 2009).[3]

This paper is rooted in the notion that a firm's production function drives its organization—that the technological interdependencies among a firm's decisions shape the organization of those who make them. As such, we expect a firm with a modular product to have a distinct organization and a firm's organization to change as its product becomes more modular. Our goal is to understand the impact of modular production on the organization of firms.

To this end, we develop a model of a single firm with two key ingredients. The first is a modular production function, which we model as a network of decisions that is partitioned into *modules*, sets of decisions that require more coordination with each other than with decisions in the other modules. The second ingredient is a communication network. Each decision is made by a different agent who observes the relevant local conditions. The communication network specifies whom each agent tells about his local conditions, after which they all make their decisions. As in Arrow (1974), each communication link comes at a cost, capturing the time and energy it takes to communicate. The organizational problem is to design an optimal communication network, trading off the efficiency of decision making with the cost of communication.[4]

The challenge in designing an optimal network is the abundance of possibilities and absence of any apparent way to order them. Our main result shows that, despite the rich set of possibilities, optimal communication networks are characterized by a simple threshold rule: each agent tells his

[1] See Peterson (2011) and Tadelis and Williamson (2013).

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

[3] Herbert Simon anticipated the rise of modular product designs in an article in 1962, in which he observed that complex systems—large firms, mechanical watches, the human body—tend to be made up of modules, groups of elements with stronger within than across group interactions (Simon 1962). The prevalence of modular structures has since been corroborated by the literature on *community detection*, which has documented them in a wide variety of complex systems, from the internet to the global air transportation network and the brain (Guimera et al. 2005, Meunier et al. 2009, and Fortunato 2010).

[4] As Kenneth Arrow put it: "*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*" Arrow (1974, p.54).

state to the other agents in his module and to all the agents in modules whose cohesion is above a threshold, where *module cohesion* captures how distinct a module is from the rest of the production network. The threshold differs across agents depending on the characteristics of their decisions and modules and the extent of uncertainty about their local conditions.

This characterization has implications for what optimal communication networks look like, what structures they exhibit. Specifically, it implies that optimal communication networks have the *threshold property*. Loosely speaking, they exhibit a common receiver ranking that orders agents by how many others tell them about their local conditions. The threshold property implies that optimal communication networks cannot take the form of a *tree* or a *matrix*, which are structures that are commonly used to describe the allocation of authority in firms. Instead, we show that under a natural condition, optimal communication networks have a *core-periphery structure*, in which modules are partitioned into an intensely communicative core and a sparsely communicative periphery. Such structures are prevalent among social networks.[5] Even if the condition is not satisfied, optimal communication networks still resemble core-periphery structures. This is so because exhibiting the threshold property implies a *generalized core-periphery structure* that differs from the standard one because of the presence of a third group of modules whose members communicate too much to be in the periphery and too little to be in the core.

Even though modular production has so far received little attention in economics, it has long been the focus of a literature in management and software engineering.[6] The hallmark of this literature is the *Mirroring Hypothesis*, which contends that organizational design mirrors product design, 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, collocation, and employment relations*" (Colfer and Baldwin (2016), p. 713). This contention leads to the prediction that the design of modular products causes firms to adopt fragmented organizations that mirror the modular structures of their products. We revisit the Mirroring Hypothesis in an application, in which we explore how an optimal communication network changes as the product becomes *more modular*. This comparative static corroborates the hypothesis by making precise why modular production leads to the fragmentation of organizations and what form this fragmentation takes.

The story of the first modular computer—IBM's *System/360*—illustrates the notion that mod-

[5] See, for instance, Borgatti and Everett (2000) and Rombach et al. (2017) and the references therein.

[6] See Thompson (1967), Conway (1968), Henderson and Clark (1990), Sanchez and Mahoney (1996), and, for a discussion of the literature, Colfer and Baldwin (2016).

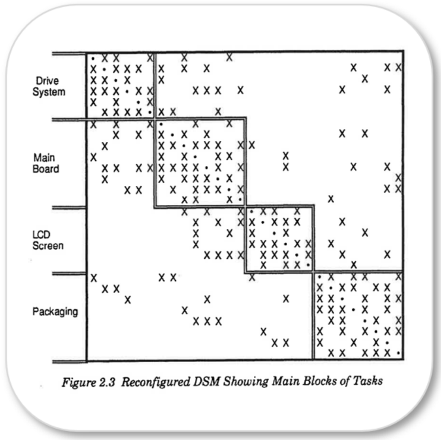

Figure 1: The *Design Structure Matrix* of a laptop in which each row and column corresponds to a task involved in designing the laptop and an "x" entry indicates a strong need for coordination between the corresponding tasks (replication of Figure 2.3 in McCord and Eppinger (1993)).

ular production impacts organization.[7] Before the System/360, computers had been tightly integrated systems of their constituent parts. As a result, a change in a single critical component required the design of an entirely new computer. This feature made it difficult to adapt computers to changes in customer preferences and led IBM to seek a computer that could be made by assembling exchangeable modules. To design the new computer, IBM changed its organization. Engineers were divided into teams, each of which was put in charge of designing a different module. Across-team communication was limited, both because the teams were scattered across the globe and because the modular structure often made it unnecessary. This fragmented organization appears to have served IBM well. The System/360 became an enormous financial success and changed how computers have been designed ever since.[8]

Our approach to modeling modular production follows a path taken in systems engineering, which characterizes products with *Design Structure Matrices* (Eppinger and Browning 2012). Figure 1 provides an example of such a matrix for a laptop. Each row and column refers to a task involved in designing the product, and the matrix entry indicates the need for coordination between

[7] This account is based on Baldwin and Clark (1997, 2000).

[8] Baldwin and Clark (1997) argue that the organizational changes still reverberate today: "*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.*"

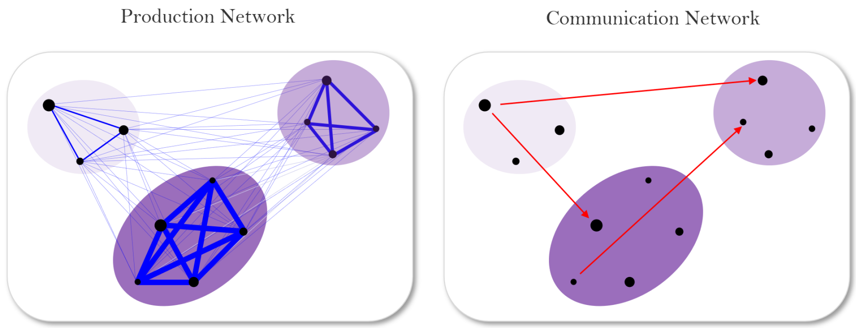

Figure 2: Left panel—The production function takes the form of a network with a *non-overlapping community structure*, where darker shading indicates more cohesive modules. Right panel—Given the production network, the principal designs optimal communication networks by deciding whom each agent should tell his state to, given that each directed link comes at an exogenous cost.

the tasks. A product is *modular* if the Design Structure Matrix is a block matrix, as it is, at least approximately, in the case of the laptop. The blocks of the matrix form the modules, such as the LCD screen. The block structure implies a greater need for coordination between tasks involved in the LCD screen than between tasks in the LCD screen and, say, those in the Main Board.

Following this approach, we model the production function as a network of decisions with a *non-overlapping community structure* (Girvan and Newman 2002), as illustrated in the left panel in Figure 2. 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 state, and the width of a link represents the importance of coordinating the two decisions it connects. Decisions are partitioned into groups of decisions that require more coordination with each other than with decisions in other groups. These *modules* are indicated by the shaded areas in the figure. This specification of the production network gives rise to an adjacency matrix with a block structure, similar to the Design Structure Matrix in Figure 1.

As we noted above, a key characteristic of the production network is the cohesions of its modules, which captures how distinct each is from the rest of the network. A module's cohesion is increasing in the number of decisions it contains and the need for coordination among them and is decreasing in the *degree of coupling*, the need for coordination between two decisions that belong to different modules. In Figure 2 a more cohesive module is indicated by darker shading.

This specification of the production network allows for a wide variety of differences in the technological structures of modular products. It allows for decisions to differ in their needs for adaptation and the degree of uncertainty about their states, for modules to differ in the number of decisions and the need for coordination between them, and for an arbitrary number of decisions and modules. The only substantive restriction is that coupling is homogeneous, that the need for coordination is the same across any two modules. We show in an extension that our results generalize naturally if we allow for heterogeneous coupling.

We embed this production network in a model of a firm's internal organization. In line with the literature on team theory, we abstract from incentive conflicts. If the agents knew all the states, their decisions would be efficient. The agents, though, only observe their own states. What they learn about the other states is determined by a communication network, such as the one illustrated in the right panel in Figure 2. Each node is the same as in the production network but the links are now directed and unweighted and indicate who tells whom his state. As we noted above, we follow Arrow (1974) in assuming that each directed link comes at a cost, which represents the time and effort it takes agents to understand each other, to learn the other's language or *code*.[9] To capture the notion that agents who belong to the same module share similar experiences, and thus use the same code, we assume that communication between them is free. In an extension we show that our results generalize readily if this assumption does not hold.

The principal's problem is to design the communication network. She takes the production network as given and builds the communication network, taking into account that each link comes at a cost. The contribution of this paper to provide an analytical solution to this problem.

To get to the solution, suppose the principal adds a directed link to an arbitrary communication network. The marginal benefit of adding such a link is the additional expected revenue generated by the agents making better decisions. 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. Communication facilitates coordination, which fosters adaptation.

A first key step to solving the principal's problem is finding that this marginal benefit is independent of what the receiver, or any other agent, knows about any other state. This feature implies that the principal's problem can be decomposed into independent subproblems. The principal can consider each agent separately and ask whom this agent should tell about his state.

[9] As Kenneth Arrow put it: "*...I am thinking of the need for having made an adequate investment of time and effort to be able to distinguish one signal from another. Learning a foreign language is an obvious example of what I have in mind. The subsequent ability to receive signals in French requires this initial investment. There are in practice many other examples of codes that have to be learned in order to receive messages; the technical vocabulary of any science is a case in point. The issue here is that others have found it economical to use one of a large number of possible coding methods, and for any individual it is necessary to make an initial investment to acquire it...It is also easier to communicate with other individuals with whom one has a common approach or a common language, literally or metaphorically. The capital accumulation of learning a code, referred to earlier, may have to be engaged in at both ends of the channel*" Arrow (1974, pp.39-42). See Cremer, Garciano, and Prat (2007) for a formal investigation of this notion.

The solution to each subproblem is driven by another property of the marginal benefit of adding a link to a communication network: the marginal benefit is larger, the more agents know the sender's state. More generally, in equilibrium, expected revenue is supermodular in the set of agents who know any given state. This property implies that if the principal benefits from adding a link to a communication network she must also benefit from adding links from the sender to all the other agents in the receiver's module.

The second key step to solving the principal's problem is finding that the same logic applies across modules: if the principal benefits from adding a link across modules, she must also benefit from adding links from the sender to any agent in a module that is more cohesive than the receiver's.

Our main result then follows readily: in an optimal communication network, each agent tells his state to the other agents in his own module and to any agent in another module whose cohesion is above a threshold. The threshold is lower, the more cohesive the sender's module is, the more important it is to adapt his decision to his state, and the more uncertainty there is about his state.

As we observed above, this result has implications for the structure of communication networks and can be applied to explore the Mirroring Hypothesis. We defer a further discussion of both until after we have presented the model and formally derived the main result.

2 Related Literature

The existing literature on modular production is largely informal and lies mostly outside of economics. It goes back to Simon (1962), who 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 software engineering, 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.

We are not aware of papers that formalize these observations and do not attempt to do so in this paper. Instead, we follow the approach of a related literature in management that takes the modular design of products as given and explores its implications for the organization of firms. As we noted earlier, a central argument in this literature is the Mirroring Hypothesis, which posits that the internal organization of firms mirrors the modular design of the products they make (see the references in Footnote 6).[10] Langlois and Robertson (1992) observed that modular production

[10] A related literature reverses the causality of the Mirroring Hypothesis and argues that the design of products mirrors the organization of the firms that designed them. In this view, a modular organization tends to design modular products. In software engineering, this view is known as *Conway's Law*, named after Melvin Conway who observed

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. We take a first step towards examining these issues through the lens of an economic model. In doing so, we take the boundaries of firms and the structure of their industries as given and examine the impact of modular production on their internal organization.

We focus, in particular, on the impact of modular production on communication structures. Such structures have long been recognized as an elemental feature of organizations (Arrow 1974). Moreover, recent empirical papers demonstrate that records of electronic communication make the patterns of such communication observable to outsiders (Yang et al. 2021, Impink, Prat, and Sadun 2021). This trend suggests that it may become possible to test predictions about communication structures and the information flows they generate, and to do so more readily than predictions about other aspects of internal organization (such as the allocation of decisions rights, which have received much attention in the theoretical literature but have proven difficult to study empirically).

Communication structures, as well as information processing, are the focus of the large and long-running literature on team theory (for an early treatment see Marschak and Radner (1972) and for more recent surveys see Garicano and Prat (2013) and Garicano and Van Zandt (2013)). A central assumption in this literature is that agents share the same goal, but cognitive constraints make it difficult for them to communicate and process information. Our focus on communication structures and cognitive constraints places us firmly in this literature.

Communication structures depend on the technological interdependencies among the decisions agents make. In many settings, this interdependency arises because decisions have to be both adapted to local conditions and coordinated with each other. March and Simon (1958) observed that this interdependency gives rise to a trade-off between adaptation and coordination that shapes the organization of firms. The contemporary organizational economics literature that explores how this trade-off affects organizations started with Dessein and Santos (2006), who explored implications for job design, and Alonso, Dessein, and Matouschek (2008) and Rantakari (2008), who examined implications for the allocation of decision rights.

We relate to a set of papers that explore how the trade-off between adaptation and coordination shapes communication structures (Calvó-Armengol and de Martí Beltran 2009, Calvó-Armengol, de Martí, and Prat 2015, Dessein, Galeotti, and Santos 2016, and Herskovic and Ramos 2020). These papers differ on various dimensions. Some focus on cognitive constraints while others allow

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).

for incentive conflicts. Some examine the centralized design of communication structures while others study their decentralized formation. And some assume that decisions need to be adapted to different states while others require them to be adapted to different signals about the same state. An assumption that is shared by all but one of these papers, though, is that the production network is complete, that the need for coordination between any two decisions is the same. This assumption rules out richer technologies, such as modular production.

The paper that allows for richer production networks, and is closest to ours, is Calvó-Armengol, de Martí, and Prat (2015). They explore an organization in which each agent adapts his decision to the local conditions about which he is privately informed. In contrast to the above papers, but like us, they allow the needs for coordination to differ across decision pairs. They do not, however, 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.

To explore the impact of modular production on communication structures, we make use of the large toolbox of network economics. The payoff functions of our agents are quadratic, and their actions are continuous and exhibit strategic complementarities. This allows us to 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), Golub and Morris (2017), and Lambert, Martini, and Ostrovsky (2018) characterized optimal decision-making for general information and network structures. We draw on their results to determine the agents' decision-making for given communication networks. Our focus, though, is not on decision-making but on the prior stage in which the principal designs the communication network.

Finally, our paper contributes to a small but growing literature that studies centralized network design.[11] In an early paper in this literature, Baccara and Bar-Isaac (2008) explored 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 were motivated by the optimal design and defense of computer networks. Even though we also explore centralized network design, we differ from these papers in both motivation and model.

[11] This literature is distinct from the large literature on endogenous network formation that started with Jackson and Wolinsky (1996) and Bala and Goyal (2000) and studies the emergence of networks from the decentralized decisions of agents. Some of the papers on communication structures we mentioned above, such as Herskovic and Ramos (2020), belong to this literature. In our model, instead, the network does not emerge endogenously from agents' communication decisions but is designed centrally by the principal.

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 $d_i \in \left[-\overline{D}, \overline{D}\right]$ that is associated with a state $\theta_i \in \left[-\overline{\theta}, \overline{\theta}\right]$, where $\mathcal{N} = \{1, \ldots, N\}$ is the set of agents, and $\overline{D}$ and $\overline{\theta}$ are large but finite scalars. Revenue 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 revenue is given by

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

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.[12] 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.

Modules. The decisions, and their associated states and agents, are partitioned into *modules*. There are M modules, and module $m \in \mathcal{M} = \{1, \ldots, M\}$ contains $n_m \geq 1$ decisions. Function $m(i)$ denotes the module 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 1, and assume that there are at least three modules, that is, $M \geq 3$.

The need for coordination between two decisions is stronger if they belong to the same module than if they belong to different ones. Specifically, the need for coordination between any two decisions d_i and d_j is given by $p_{ij} = p \geq 0$ if they belong to different modules and, abusing notation slightly, it is given by $p_{ij} = p_m \geq p$ if they belong to the same module m. The parameter p,

[12] A special case of this formulation is 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 revenue as

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

where the last term is a constant.

therefore, captures the *degree of coupling*—the need for coordination across modules—while the parameter p_m captures the need for coordination within module 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.

Before the states are drawn, the principal can place communication links from any agent to any others. A link from one agent to another is free if they belong to the same module but costs $\gamma > 0$ if they belong to different ones. This assumption captures the notion that it takes resources for one agent to tell his state to another, especially if they do not share the same 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 one if agent i tells agent j 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. The communication network, and all other information except for the agents' states, are common knowledge.

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

$$\max_{\boldsymbol{C}} \mathrm{E}[r(d_1, \ldots, d_N) \,|\, \boldsymbol{C}] - \gamma \sum_{i=1}^{N} \sum_{j=1}^{N} m_{ij} c_{ij}, \tag{2}$$

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

Timing. After the principal designs the communication network, agents learn their states and tell them to the other agents as specified by the network. Next, agents simultaneously make their decisions, after which the game ends.

We discuss key assumptions, such as homogeneous coupling, the agents' inability to re-transmit information, and the absence of incentive conflicts in Section 7, after we solve the model in the next section, derive implications in Section 5, and apply it to the Mirroring Hypothesis in Section 6.

4 Solving the Model

We start by determining Bayes-Nash 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 decomposing it into independent subproblems. Finally, we characterize the solution to the principal's problem by solving these subproblems.

4.1 Decision-Making

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

$$\max_{d_i} \mathrm{E}\left[r\left(d_1, \ldots, d_N\right) \middle| \boldsymbol{C}_{(i)}\right] \text{ for all } i \in \mathcal{N}, \tag{3}$$

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

$$d_i = a_i \theta_i + \sum_{j=1}^{N} p_{ij} \mathrm{E}\left[d_j \middle| \boldsymbol{C}_{(i)}\right]. \tag{4}$$

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 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 $(\mathrm{diag}\boldsymbol{C}_j)\, \boldsymbol{P}\, (\mathrm{diag}\boldsymbol{C}_j)$ is the subgraph of the production network that consists of the nodes whose agents know state θ_j, and all the links between them. We can then state the following lemma.

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

$$d_i^* = \sum_{j=1}^{N} a_j \omega_{ij}\left(\boldsymbol{C}_j\right) \theta_j \textit{ for all } i \in \mathcal{N}, \tag{5}$$

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

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 $\left(\boldsymbol{I} - (\mathrm{diag}\boldsymbol{C}_j)\, \boldsymbol{P}\, (\mathrm{diag}\boldsymbol{C}_j)\right)^{-1}$. This latter object has a natural interpretation in terms of walks on the production network: it is 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.[13] If agent i does not know θ_s, for instance, d_i is not part of this subgraph, and so $\omega_{is}\left(\boldsymbol{C}_s\right) = 0$. 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: 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.

[13] 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.

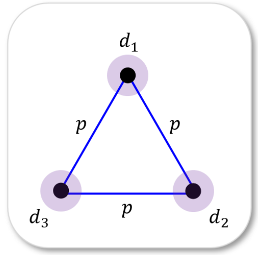

Figure 3: Production network $\boldsymbol{P}^e$.

The part of the equilibrium decision rules that will turn out to be important for the design of communication networks is the weight each agent's decision puts on his own state. To get an intuition for this weight, consider the production network $\boldsymbol{P}^e$ in Figure 3, where the superscript stands for *example*. Suppose first 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

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

Suppose instead 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 puts 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 he puts on θ_1 increases to

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

Communication enables coordination, which, in turn, facilitates adaptation. The extent to which it does so is captured by the *coordination multiplier* $\omega_{11}(\boldsymbol{C}_1)$.[14]

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 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

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

where the last term in brackets captures the supermodularity.

[14] The coordination multiplier is related to the notion of *cycle centrality* in Talamàs and Tamuz (2017).

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

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

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

4.2 Simplifying the Principal's Problem

The principal's problem is to design a communication network that maximizes expected profits. It is useful to start by rewriting revenue (1) as

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

Substituting in the best-response functions (4), this expression simplifies to

$$r(d_1^*, \ldots, d_N^*) = \sum_{i=1}^{N} a_i d_i^* \theta_i + \sum_{i=1}^{N}\sum_{j=1}^{N} p_{ij} d_i^* \left(d_j^* - \mathrm{E}\left[d_j^* \,\middle|\, \boldsymbol{C}_{(i)} \right] \right). \tag{6}$$

Notice that by the law of iterated expectations the second term on the right-hand side is zero in expectation, which delivers the following result.

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

$$R(\boldsymbol{C}) \equiv \mathrm{E}\left[r(d_1^*, \ldots, d_N^*) \right] = \sum_{i=1}^{N} a_i \mathrm{Cov}(d_i^*, \theta_i), \tag{7}$$

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

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

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

The term $a_i^2\sigma_i^2$ is the revenue agent i is expected to generate if he does not tell his state to any other agent and, thus, adapts his decision to his state autonomously. We refer to this term as the *value of autonomous adaptation* of decision d_i. The coordination multiplier captures how much more revenue agent i is expected to generate when he adapts his decision more closely to his state because other agents know his state.

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 $\omega_{ii}(\boldsymbol{C}_i)$ and thus the weight $a_i\omega_{ii}(\boldsymbol{C}_i)$ he puts on his state. As a result, it also increases the expected revenue $a_i^2\sigma_i^2\omega_{ii}(\boldsymbol{C}_i)$ he generates. In contrast, agent i, or any other agent, learning any other state does not affect $\omega_{ii}(\boldsymbol{C}_i)$ and thus leaves the weight agent i puts on his own 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 separately 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 the following.

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

$$\max_{\boldsymbol{C}_i} R_i(\boldsymbol{C}_i) - \gamma\sum_{j=1}^{N} m_{ij}c_{ij}, \tag{8}$$

for each $i \in \mathcal{N}$, where m_{ij} is a dummy variable that is equal to one if agents i and j belong to different modules.

This separability result greatly facilitates the principal's quest for an optimal communication network. We can further simplify the problem by recalling that agent i's coordination multiplier $\omega_{ii}(\boldsymbol{C}_i)$ 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. The principal's problem, therefore, reduces to asking which modules each agent should tell about his state.

Finally, supermodularity of $\omega_{ii}(\cdot)$, together with the linearity of communication costs, implies 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 separately and asking whom he should tell about his state. The first step in answering this question is to express the agents' expected revenues in terms of the model's primitives. To economize on notation, and without loss, we start by focusing on agent 1.

Suppose agent 1 tells his state to all the agents in an arbitrary set of modules that includes his own module 1. Since the naming of modules is immaterial, there is no loss in denoting the modules in this set by $\{1, \ldots, \ell\}$, where $\ell \in \mathcal{M}$. We define $\boldsymbol{C}_1(\ell)$ as the row of the communication matrix that specifies the agents who belong to modules $1, \ldots, \ell$ 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 the agents in modules* $1, \ldots, \ell$, *where* $\ell \in \mathcal{M}$, *and to none of the agents in the other modules* $\mathcal{M} \backslash \{1, \ldots, \ell\}$. *Agent* 1*'s expected revenue is then given by*

$$R_1(\boldsymbol{C}_1(\ell)) = a_1^2 \sigma_1^2 \left(\frac{1 + (p_1 - p) x_1}{1 + p_1} + \frac{p x_1^2}{1 - p \sum_{m=1}^{\ell} n_m x_m} \right), \tag{9}$$

where

$$x_m = \frac{1}{1 + p - (n_m - 1)(p_m - p)} \text{ for } m \in \mathcal{M}.$$

The object x_m in the lemma is the *module cohesion* of module $m \in \mathcal{M}$, which captures how distinct the module is from the rest of the production network. To see this interpretation, note that the last term in the denominator is the *excess need for coordination* of any decision d_i in the module, the difference between its total need for coordination $\sum_{j=1}^{N} p_{ij} = (n_m - 1) p_m + (N - n_m) p$ and $(N-1)p$, the value $\sum_{j=1}^{N} p_{ij}$ would take if the need for coordination within module m were the same as that across the modules. A module is, therefore, more cohesive than another module if and only if its decisions have a higher excess need for coordination.[15]

We already know that agent 1's expected revenue is the product of $a_1^2\sigma_1^2$—the value of autonomous adaptation of his decision—and the coordination multiplier $\omega_{11}(\boldsymbol{C}_1(\ell))$. The lemma shows that the coordination multiplier takes a simple form. Notice, in particular, that the coordination multiplier depends on the characteristics of modules $2, \ldots, \ell$ only through the sum of their scaled cohesions $n_2 x_2 + \cdots + n_\ell x_\ell$ and that it is convex in this sum. This convexity reflects the supermodularity of the coordination multiplier we noticed earlier. It is important here because it implies that if it is profitable to expand the set of informed modules by one additional module, it must also be profitable to expand it further by adding a second module, provided that the second is no less cohesive than the first. The answer to whom agent 1, or any other agent, should tell about his state follows directly from this claim.

[15] Our definition of module cohesion is close to the definition of cohesion in Morris (2000). Applied to our setting his, like ours, is increasing in n_m and p_m and decreasing in p.

PROPOSITION 2. *There exist thresholds $\lambda_i \geq 0$ such that it is optimal for agent $i \in \mathcal{N}$ to tell his state to a different agent $j \in \mathcal{N}$ if and only if:*

(i.) agent j belongs to the same module $m(j) = m(i)$ or

(ii.) agent j belongs to a different module $m(j) \neq m(i)$ with cohesion $x_{m(j)} \geq \lambda_i$.

The threshold λ_i is increasing in γ, decreasing in $a_i^2\sigma_i^2$, p_m, and n_m for any $m \in \mathcal{M}$, and independent of $a_k^2\sigma_k^2$ for any $k \in \mathcal{N} \setminus \{i\}$.

The key result in the proposition is that, apart from the agents in his own module, an agent should tell his state to the agents in the most cohesive modules. To see the intuition for why cohesion matters, suppose agent 1 tells his state to the agents in one additional module, say module 2. Since the agents in module 2 care about coordinating their decisions with agent 1's, this induces them to put some weight on θ_1. And since agent 1 also cares about coordinating his decision with theirs, this, in turn, allows him to adapt his decision more aggressively to his state. Communication facilitates coordination, which fosters adaptation, as we noted before.

The agents in module 2, though, do not only care about coordinating their decisions with agent 1's, they also care about coordinating them with each others'. This is where the cohesion of module 2 comes into play. The more cohesive the module is—the higher the excess need for coordination of its decisions is—the *more* its agents care about coordinating their decisions with each other and, thus, the *more* weight their decisions put on θ_1. Cohesion, therefore, matters because it strengthens the link between communication and coordination which, ultimately, leads to more adaptation.

To get an intuition for the comparative statics in the proposition, consider Figure 4, where we again focus on agent 1. There are five modules, with modules $2, \ldots, 5$ labeled in decreasing order of their cohesion: $x_2 \geq x_3 \geq x_4 \geq x_5$. The red line is a continuous representation of communication costs $\gamma \sum_{m=2}^{\ell} n_m$ and the blue curve is the piecewise-linear extension of expected revenue $R_1(\boldsymbol{C}_1(\ell))$, which we denote by $\overline{R}_1(\boldsymbol{C}_1(\ell))$. The changing curvature of expected revenue $\overline{R}_1(\boldsymbol{C}_1(\ell))$ reflects the countervailing economic forces at work. Supermodularity pushes towards convexity while the modular structure of the production function pushes towards concavity.

Now consider the effect of changes in the parameters on the cost and benefit curves. The slope of each line segment in the benefit 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. This per node marginal benefit is larger, the higher the value of autonomous adaptation for agent 1's decision is. And it is larger, the more agents know θ_1 and the higher the need for coordination among them is. An increase in $a_1^2\sigma_1^2$, n_m, or p_m, therefore, steepens the benefit curve, which favors telling agents in more modules about θ_1. In contrast, an increase in the communication costs γ

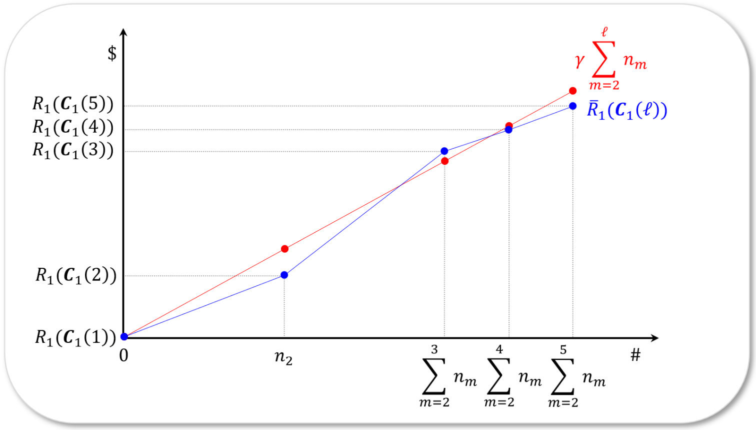

Figure 4: Determining optimal communication networks for agent 1 (drawn for parameter values $n_1 = n_2 = n_3 = 5$, $n_4 = n_5 = 2$, $p_1 = p_2 = p_3 = 0.2$, $p_4 = p_5 = 0.1$, $p = 0.01$, and $a_1\sigma_1 = 1$).

steepens the cost curve, which favors telling agents in fewer modules about θ_1.

The characterization of the optimal communication networks in Proposition 2 is our main result. To derive it, we first showed that the principal's problem can be decomposed into independent subproblems. While the subproblems are independent, though, their solutions are related in a way that has implications for what types of structures optimal communication networks exhibit. We explore these implications next.

5 Implications for Network Structures

We now turn to the implications of the characterization result in Proposition 2 for the structure of optimal communication networks. Our starting point is that optimal communication networks have the *threshold property*. Once we have established this result, we show that the class of networks that have this property excludes a number of well-known structures, and includes others.

Intuitively, a communication network has the threshold property if it exhibits a *common* receiver ranking—if agents can be ranked by whether they are told about a state in a different module, and the *same* ranking determines the order in which they are told about any other state that belongs to a different module. Formally, we have the following.

DEFINITION. *A communication network has the "threshold property" if there exist sender thresholds $\{s_1, \ldots, s_N\}$ and receiver scores $\{r_1, \ldots, r_N\}$ such that for any two agents $i, j \in \mathcal{N}$ who belong to different modules, agent i tells agent j his state if and only if $r_j \geq s_i$.*

In line with the above intuition, the definition requires a common receiver ranking. Suppose

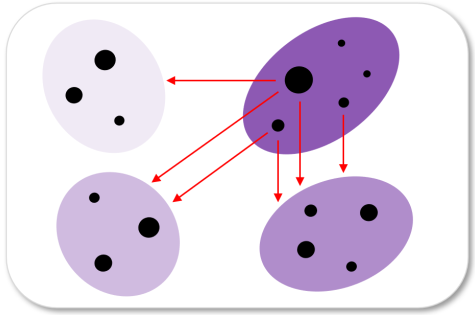

Figure 5: Across-module communication in an optimal communication network exhibiting the *threshold property*, where an arrow from an agent to a module indicates communication links from the agent to all the agents in the module, and darker shading indicates more cohesive modules.

agents i and j belong to different modules, and agent i's receiver score is larger than agent j's, $r_i > r_j$. Agent i then has a higher receiver rank than agent j does, he is told about any state in any module $m(k) \in \mathcal{M} \setminus \{m(i), m(j)\}$ that agent j is told about, and possibly others.

Optimal communication networks exhibit a common receiver ranking. Suppose agent i's module is more cohesive than agent j's, $x_{m(i)} > x_{m(j)}$. The characterization result in Proposition 2 then implies that, in an optimal communication network, agent i is told about any state in any module $m(k) \in \mathcal{M} \setminus \{m(i), m(j)\}$ that agent j is told about, and possibly others. We can, therefore, obtain the receiver scores the definition calls for by setting $r_i = x_{m(i)}$, and the sender thresholds by setting $s_i = \lambda_i$, for all $i \in \mathcal{N}$. Doing so delivers the result, which we also illustrate in Figure 5.

COROLLARY 2. *Optimal communication networks have the threshold property.*

The class of networks that exhibit the threshold property excludes a number of well-known structures, and includes others. To describe what it excludes, we adapt Cloteaux et al. (2014) and show that the threshold property rules out certain forbidden subgraphs. This then allows us to rule out any structure that contains one or more of these subgraphs. The relevant subgraphs are *two-switches* and *directed three-cycles*, which we illustrate in Figure 6 and define as follows.

DEFINITION. *A communication network contains a "two-switch" if there are four agents i, j, k, $l \in \mathcal{N}$ who belong to different modules such that agent i tells his state to k but not to l and agent j tells his state to l but not to k.*

A communication network contains a "directed three-cycle" if there are three agents $i, j, k \in \mathcal{N}$ who belong to different modules such that agent i tells his state to j, but not the reverse, agent j his state to k, but not the reverse, and agent k tells his state to i, but not the reverse.

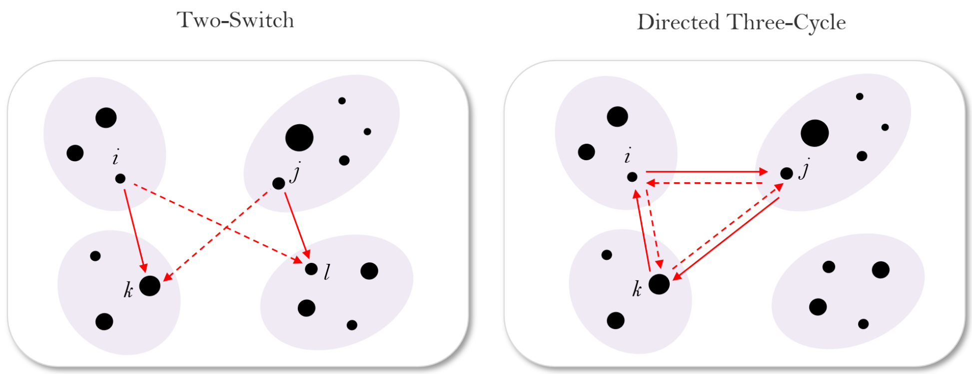

Figure 6: Communication networks exhibiting two-switches and directed three-cycles, where a dashed arrow indicates the absence of a communication link.

The next lemma shows that the threshold property rules out these subgraphs.

LEMMA 4. *A communication network with the threshold property contains no two-switches or directed three-cycles.*

To see why two-switches cannot be part of an optimal communication network, consider agents i, j, k, and l from the definition of two-switches and their illustration in Figure 6. If it is optimal for agent i to tell his state to agent k but not to agent l, it must be that module $m(k)$ is more cohesive than module $m(l)$. If module $m(k)$ is more cohesive than module $m(l)$, though, it cannot be optimal for agent j to tell his state to agent l but not to agent k. The intuition for why three-cycles cannot be part of an optimal communication network is similar.

The lemma is useful because it rules out the many well-known network structures that do contain two-switches or three-cycles (or both). Two structures are of particular relevance to us: trees and matrices. These structures—which we illustrate in Figure 7 and define below—are commonly used to describe the allocation of decision rights in firms. As the examples in the figure suggest, though, they all contain at least one two-switch. It then follows from the lemma, together with Corollary 2, that any optimal communication structure must be distinct from these well-known ones.

To state the result, we first define the two structures as follows.

DEFINITION. *A communication network is a "tree" if there is a partition of the set of modules into hierarchical levels* $1,\ldots,H$ *with* $H \geq 3$ *such that there is one module in level* 1 *and at least two modules in level* $h \in \{2,\ldots,H\}$. *Each module in level* $h \in \{2,\ldots,H\}$ *is associated with a unique "predecessor module" in level* $h-1$. *Any two agents* i,$j \in \mathcal{N}$ *who belong to different modules tell each other their states if and only if* $m(i)$ *is a predecessor of* $m(j)$ *or vice versa.*

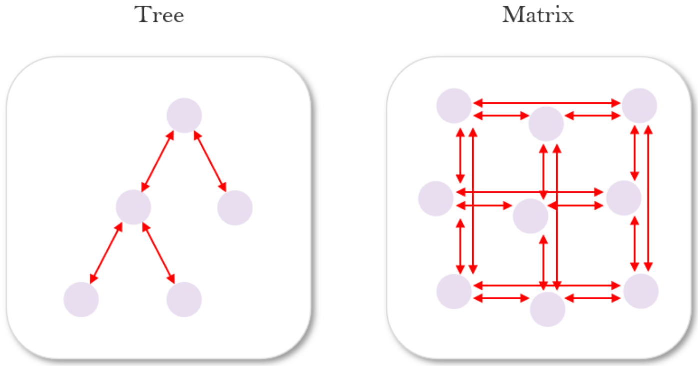

Figure 7: Illustration of trees and matrices, where a shaded area indicates a module and an arrow indicates communication links from all agents in one module to all the agents in the other.

A communication network is a "matrix" if there are two distinct partitions of the set of modules into horizontal and vertical teams, with at least two horizontal teams containing at least two modules, and such that any two modules in the same horizontal team are in different vertical teams. Any two agents $i,j \in \mathcal{N}$ who belong to different modules tell each other their states if and only if their modules belong to the same horizontal or vertical team.

We then have the following.

COROLLARY 3. *A communication network is not optimal if it is a tree or matrix.*

The fact that the threshold property rules out well-known network structures raises the question of what structures it is consistent with. We show next that under a natural condition, the threshold property gives rise to a *core-periphery structure* with intense communication in the core, no across-module communication within the periphery, and sparse communication between the core and the periphery. As we noted in the introduction, such structures are prevalent among social networks.

DEFINITION. *A communication network has a "core-periphery structure" if the set of agents 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 (ii.) an agent in the periphery does not tell his state to other agents in the periphery who do not belong to his module.*

The next proposition shows that optimal communication networks have a core-periphery structure if the agents who have a low sender threshold also have a high receiver rank. Loosely speaking, agents who talk a lot then also hear a lot, and thus form the core, while those who hear little also talk little and find themselves in the periphery.

PROPOSITION 3. *Suppose that for any agents $i, j \in \mathcal{N}$, $\lambda_i \leq \lambda_j$ if and only if $x_{m(i)} \geq x_{m(j)}$. Any optimal communication network then has a core-periphery structure in which the agents who belong to the most cohesive modules form the core.*

The condition in the proposition is satisfied if decisions that belong to more cohesive modules have sufficiently higher values of autonomous adaptation $a_i^2\sigma_i^2$ than those that belong to less cohesive ones. This follows from the results in Proposition 2 that an increase in $a_i^2\sigma_i^2$ reduces λ_i but leaves λ_j, and cohesions $x_{m(i)}$ and $x_{m(j)}$, unchanged. As such, the condition in the proposition is satisfied, for instance, if modules can be ranked by their *importance,* with more important modules being larger and their decisions requiring both more adaptation and more coordination.

Even if the condition in the proposition does not hold, the structures of optimal communication networks are akin to core-periphery structures. In particular, we show in Lemma 5 in the appendix that exhibiting the threshold property implies a *generalized core-periphery structure*: modules can be partitioned into a core, a periphery, and what we refer to as a *suburban periphery.* Suburban periphery modules are involved in too much communication to the periphery and too little communication from the core to belong to either.

A feature of the implications we derived in this section is that they do not depend on the specifics of the production network (such as the needs for adaptation and coordination and the degrees of uncertainty) beyond its modular structure. They can be tested, at least in principle, with information about what modules decisions belong to, which is contained in Design Structure Matrices, and data about whom employees communicate with, such as those in Yang et al. (2021) and Impink, Prat, and Sadun (2021). In the next section, we turn to implications of the characterization that would require more information to test empirically. These implications, though, allow us to speak to the central claim in the literature on modular production, the Mirroring Hypothesis.

6 Application to the Mirroring Hypothesis

Next we apply the characterization result to the Mirroring Hypothesis, the notion that the design of modular products causes firms to adopt fragmented organizations that mirror the modular structures of their products.

To explore the Mirroring Hypothesis through the lens of our model, we need to be precise about terms that are not always formally defined in the literature. We start with the notion that a communication network is *admissible.*

DEFINITION. *A production network* $\boldsymbol{P}$ *"admits" a communication network* $\boldsymbol{C}$ *if there exists a communication cost* $\gamma \in [0, \infty)$ *for which* $\boldsymbol{C}$ *is a solution to the principal's problem* (2)*.*

Two communication networks are always admissible: if communication costs are sufficiently low, it must be optimal for all agents to tell their states to all the others, and if they are sufficiently high, it must be optimal for all agents to tell their states to the others in their own modules but to no one else. In these communication networks any across-module communication is firm-wide. We are interested in when there are other communication networks that are admissible, ones in which there is at least one agent who tells his state to some agents in other modules but not to all of them. We call such communication networks *fragmented*.

DEFINITION. *A communication network* $\boldsymbol{C}$ *is "fragmented" if there is an agent* $i \in \mathcal{N}$ *who tells his state to one or more agents in other modules* $\mathcal{M} \backslash \{m(i)\}$ *but not to all of them.*

The definitions of admissibility and fragmentation extend readily from communication network $\boldsymbol{C}$ to its row vector $\boldsymbol{C}_i$. Row vector $\boldsymbol{C}_i$ is admissible if there exists a communication cost for which it is a solution to the principal's subproblem (8), and it is fragmented if agent i tells his state to some agents in other modules but not to all of them. In these terms, communication network $\boldsymbol{C}$ is admissible only if all its row vectors $\boldsymbol{C}_1, \ldots, \boldsymbol{C}_N$ are, and it is fragmented if at least one of its row vectors $\boldsymbol{C}_i$ is.

The admissibility and fragmentation of $\boldsymbol{C}_i$ is determined by the properties of agent i's expected revenue. To illustrate this point, recall the example in Figure 4, which plots the piecewise-linear extension of agent 1's expected revenue $\overline{R}_1(\boldsymbol{C}_1(\ell))$, where $\boldsymbol{C}_1(\ell)$ is the row vector that has agent 1 tell his state to the agents in modules $1, \ldots, \ell$. It is admissible if and only if $\overline{R}_1(\boldsymbol{C}_1(\ell))$ lies on its concave closure. There are communication costs for which agent 1 should tell his state only to the agents in module 1, only to those in modules 1 to 3, or to the agents in all five modules. In contrast, there are no communication costs for which agent 1 should tell his state only to the agents in modules 1 and 2 or only to those in modules 1 to 4. These row vectors are not on the concave closure of $\overline{R}_1(\boldsymbol{C}_1(\ell))$ and are, thus, not admissible. In this example, $\boldsymbol{C}_1(3)$ is, therefore, the only fragmented row vector that is admissible, that is optimal for some communication costs.

The economic force that pushes against fragmentation is the supermodularity of the coordination multiplier, and thus expected revenue, while the property that pushes in its favor is the modular structure of the production network. The first result in this section shows that if a production network is only *weakly modular*—if the within-module needs for coordination are not much stronger than the across-module ones—supermodularity dominates, and fragmented organizations are not admissible. To state this result, we first define weak modularity as follows.

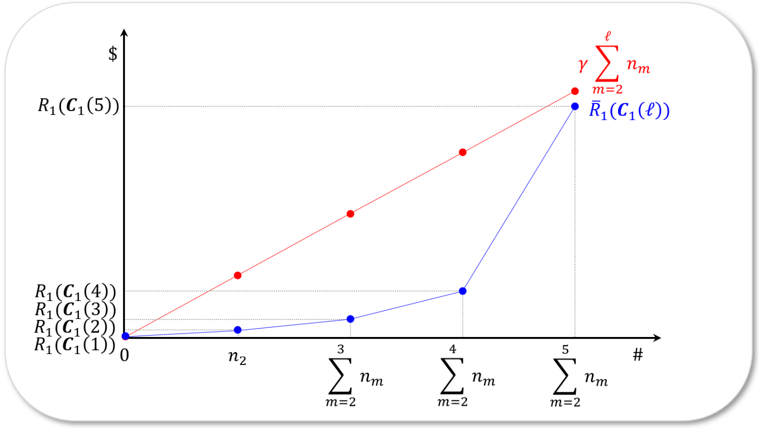

Figure 8: Determining optimal communication networks for agent 1 when the production network is *weakly modular* (drawn for parameter values $n_1 = n_2 = n_3 = n_4 = n_5 = 2$, $p_1 = 0.15$, $p_2 = 0.14$, $p_3 = 0.13$, $p_4 = 0.12$, $p_5 = 0.11$, $p = 0.1$, and $a_1\sigma_1 = 1$).

DEFINITION. *Production network* $\boldsymbol{P}$ *is "weakly modular" if*

$$p_m - p < \frac{2p\,(1+p)}{(1+3p)\,(n_m-1)} \text{ for all } m \in \mathcal{M}.$$

We then have the following.

PROPOSITION 4. *If production network* $\boldsymbol{P}$ *is weakly modular, it does not admit any fragmented communication networks.*

Intuitively, if the production network is only weakly modular, its modules cannot be very pronounced. Their cohesions must be small and, thus, similar to each other. The result then follows from supermodularity: if the modules are similar, it may be optimal for an agent to tell his state to the agents in all the other modules or to none of them. Mixing and matching—telling the state to the agents in some modules but not to those in others—cannot be optimal.

We illustrate the proposition in Figure 8, which plots the piecewise linear extension of agent 1's expected revenue $\overline{R}_1(\boldsymbol{C}_1(\ell))$ for a weakly modular production network. Because the production network is weakly modular, $\overline{R}_1(\boldsymbol{C}_1(\ell))$ is convex, which rules out fragmented row vectors from being on its concave closure and thus from being admissible.

We can now turn to how the communication network changes if the production network becomes *more modular*—if the within-module needs for coordination become stronger relative to the across-

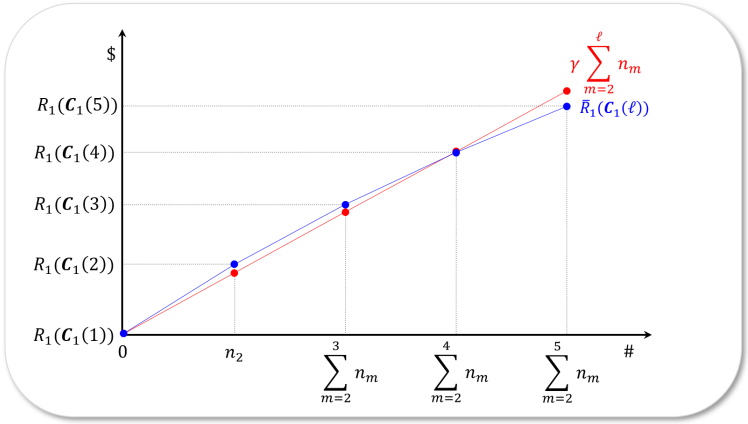

Figure 9: Determining optimal communication networks for agent 1 after a *weight-neutral increase in modularity* of size 0.99 of the production network specified in the caption of Figure 8.

module ones. Recall that the total need for coordination of decision d_i is given by

$$\sum_{j=1}^{N} p_{ij} = \left(n_{m(i)} - 1\right) p_{m(i)} + \left(N - n_{m(i)}\right) p \text{ for any } i \in \mathcal{N}. \tag{10}$$

To avoid conflating the effect of an increase in modularity with those due to changes in the total needs for coordination, we focus on increases that are *weight-neutral*, ones that do not change (10).

DEFINITION. *A "weight-neutral increase in modularity of size μ" of production network $\boldsymbol{P}$ reduces the degree of coupling from p to $p(\mu) \equiv (1-\mu)p$ and increases the within-module needs for coordination from p_m to*

$$p_m(\mu) \equiv p_m + \mu p \frac{N - n_m}{n_m - 1} \textit{ for all } m \in \mathcal{M},$$

where $\mu \in (0,1)$. We denote the resulting production network by $\boldsymbol{P}(\mu)$.

The example in Figure 9, for instance, is a weight-neutral increase in modularity of size 0.99 of the production function in Figure 8. The figure illustrates the next result: after a sufficiently large increase in modularity, some fragmented row vectors are on the concave closure of agent 1's expected revenue $\overline{R}_1(\boldsymbol{C}_1(\ell))$ and thus admissible.

PROPOSITION 5. *If production network $\boldsymbol{P}$ is weakly modular, and the modules differ in their cohesions, there exists a $\overline{\mu} \in (0,1)$ such that $\boldsymbol{P}(\mu)$ admits a fragmented communication network for all $\mu \in [\overline{\mu}, 1)$.*

The intuition follows from the effect of modularity on module cohesion. As the production network becomes more modular, its modules become more cohesive. Because the increase in cohesion is larger the more cohesive a module is, this causes the module cohesions to diverge. At some point, the modules have such different cohesions that the modular structure of the production network dominates the supermodularity of expected revenue. There then exist communication costs for which it is optimal for an agent to tell his state to agents in some other modules—the most cohesive ones—but not to those in others. This conclusion does not hold, though, if the modules all have the same cohesion before the increase in modularity. If they did have the same cohesion before the increase, they would still have the same afterwards. Supermodularity would then keep fragmented communications inadmissible, no matter how much modularity increased.

Together Propositions 4 and 5 corroborate the Mirroring Hypothesis. They provide a precise argument for why modularity is associated with fragmentation and show what forms it can take. An increase in modularity is associated with fragmentation because it makes modules more distinct from each other and does so on the dimension that matters—their cohesions. This change in the production network can lead to the removal of communication links to the least cohesive modules, the addition of links to the most cohesive ones, or both. Other ways in which modularity may fragment the communication network are not consistent with the model.

7 Relaxing Assumptions

Having solved the model and explored its implications, we return to the model assumptions in Section 3 to discuss how our results change as we relax them.

Costly within-module communication. The assumption that there are no costs to communicating within modules captures the notion that communication takes less time and effort if agents share similar experiences. The characterization of optimal communication networks in Proposition 2 extends readily to an alternative specification in which communication within modules comes at a cost, just as communication across modules does. We examine this alternative in Appendix B.

Imperfect communication. Our model follows Calvó-Armengol and de Martí Beltran (2009) in assuming that communication is perfectly informative. Suppose, instead, that when agent i tells agent j his state, agent j learns state 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 amounts to re-scaling the communication costs: the principal's problem is still given by Proposition 1 with communication costs γ replaced by γ/q.

Correlated states. An assumption we share with Calvó-Armengol, de Martí, and Prat (2015) is that states are independent. If, instead, 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 whom each agent should tell about his state would be interdependent. In such a setting, expected revenue $R(\boldsymbol{C})$ could still be written as the sum of the expected revenue generated by each agent, $R_i(\cdot)$, that is, Lemma 2 would still hold. The expected revenue generated by each agent, though, would depend on the entire communication network $\boldsymbol{C}$, causing the principal's subproblems in Proposition 1 to become interdependent.

Re-transmission of information. Next we explore the assumption 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). This assumption captures the idea that, even though we model each state as simply a number, it refers to a complex set of conditions and circumstances that only the associated agent can describe appropriately.

One way to relax this assumption is to allow agents to *synthesize* information, to combine their knowledge of their own state with information they receive about other states. Suppose, for instance, that after learning state θ_i, agent j can communicate a summary statistic of (θ_i, θ_j) to agent k and do so at cost γ. A communication link from agent i to j then does not only affect what agent j knows about θ_i but also what he can tell others. As a result, the principal's subproblems in Proposition 1 become interdependent, and the separability result no longer holds.

Another way to allow for re-transmission is to suppose that after being told state θ_i, agent j must incur cost γ to tell θ_i to agent k. Agents can tell others the states they have been told about and they can do so at the same cost at which they can tell them about their own states. In this case, any equilibrium with re-transmission is payoff equivalent to one without. The separability result in Proposition 1 continues to hold, as does the characterization result in Proposition 2.

Heterogeneous coupling. The assumption that coupling is homogenous applies to settings, such as the laptop in Figure 1, in which the needs for coordination across any two modules are (roughly) the same. Even though such settings are common, there are others in which some modules need to be coordinated more closely with each other than with other modules. The law and business schools of a university, for instance, may require more coordination with each other than with the schools of engineering and natural sciences, and vice versa.

To explore such heterogeneous coupling, suppose we partition the set of modules into clusters that differ in their degrees of coupling, such as in the example in the left panel in Figure 10. Suppose, in particular, that each node i belongs to a module $m(i)$, and each module m belongs

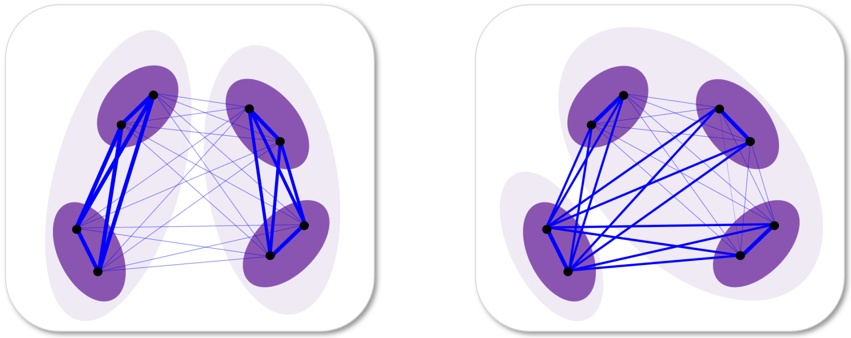

Figure 10: Left panel—Production network with *heterogeneous coupling*, where the dark shaded areas indicate modules and the lighter ones indicate clusters of modules. Right panel—Production network with an *interface module*.

to a cluster $k(m) \in \mathcal{K}$, where $\mathcal{K} = \{1, \ldots, K\}$. As in the main model, the need for coordination between any two decisions d_i and d_j is given by $p_{ij} = p_m \geq 0$ if they belong to the same module m. In contrast to the main model, however, the need for coordination between the two decisions need no longer be given by p if they belong to different modules. Instead, it is given by $p_{ij} = p^k \geq 0$ if their modules belong to the same cluster k, and it is given by p only if their modules belong to different clusters. The new parameters $p^1, \ldots, p^K$, therefore, capture the degrees of coupling in the different clusters. The rest of the model is as in Section 3.

In Appendix D we show that, even though the computations get considerably more involved, we can still derive a closed-form expression for the agents' expected revenue in terms of the model primitives, such as the one in Lemma 3. This result, in turn, allows us to show that the characterization result in Proposition 2 generalizes as follows.

PROPOSITION 6. *There exist thresholds* $\lambda_i^k \geq 0$ *such that it is optimal for agent* $i \in \mathcal{N}$ *to tell his state to a different agent* $j \in \mathcal{N}$ *who belongs to cluster* $k \in \mathcal{K}$ *if and only if:*

(i.) agent j *belongs to the same module* $m(j) = m(i)$ *or*

(ii.) agent j *belongs to a different module* $m(j) \neq m(i)$ *with cohesion* $x_{m(j)} \geq \lambda_i^k$.

Optimal communication is, therefore, still determined by threshold rules on module cohesions. Apart from the agents in his own module, each agent tells his state to the agents in modules whose cohesion is sufficiently high. The only difference is that how high cohesion needs to be depends on the characteristics of the cluster that the receiving agent belongs to.

Interfaces. A feature of modular production that is important in some applications is the presence of an *interface module*—a module that all other modules have to be tightly coordinated with. To return to the university example, the various schools of a university may all need to be coordinated more closely with the office of the provost than with each other. The extension to heterogeneous coupling allows us to incorporate this feature.

Notice, in particular, that the extension allows for the across-cluster degree of coupling p to be higher than the within-cluster degrees of coupling $p^1, \ldots, p^K$. We can, therefore, specify a production network, such as the one in the right panel in Figure 10, in which the set of modules is divided into one cluster with a single module and another cluster with all the other ones. The across-cluster degree of coupling then captures how tightly the modules in the multi-module cluster must be coordinated with the single module, and the within-module degree of coupling captures how well they must be coordinated with each other. Setting the across-module degree of coupling higher than the within-module one turns the single module into an interface module.

Because this specification is a special case of the extension with heterogeneous coupling, Proposition 6 still applies. Even with an interface module, optimal communication networks are determined by threshold rules on module cohesion.

General production networks. Even though we focus on modular production, it is instructive to examine how our results change if we allow for general production networks. To this end, suppose 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 proofs of Lemma 1, Lemma 2, and Proposition 1 allow for such production networks, so the separability result continues to hold. As such, the principal can still determine an optimal communication network by considering each agent separately. Moreover, we show in Proposition 7 in Appendix D that 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 uses the presence of modules. Optimal communication networks can now take many forms and need not exhibit the threshold property. The supermodularity of the principal's objective, though, ensures that comparative statics are still monotone. As in the main model, 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.

Incentive conflicts. Lastly, we explore how the results change if we depart from our team theoretic approach and allow for incentive conflicts. To this end, suppose that each agent only internalizes a fraction $\rho \in [0, 1]$ of the needs for coordination, that he acts as if the production network were given by $\rho\boldsymbol{P}$ rather than $\boldsymbol{P}$ (e.g., because agents put more weight on their own revenue, as in Athey and

Roberts (2001) and Alonso, Dessein, and Matouschek (2008)). The agents' decisions then solve

$$\max_{d_i} \mathrm{E}\left[\sum_{i=1}^{N}\left[-d_i^2+2a_id_i\theta_i+\rho\sum_{j=1}^{N}p_{ij}d_id_j\right]\middle|\, \boldsymbol{C}_{(i)}\right] \text{ for all } i\in\mathcal{N},$$

where the only difference with the decision problem in (3) is the ρ in front of the third term in the inner brackets. The rest of the model is as in Section 3.

We study this extension in Appendix E. In the presence of incentive conflicts, Proposition 8 shows that the principal's objective function is the same as in (8), except for an additional term equal to the weighted sum of the covariances between each decision pair. This term 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 poses is that the extent to which two decisions co-vary with each other depends on which states both decision makers know about. As a result, the principal can no longer consider each agent separately 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.

8 Conclusions

The structure of technology drives the organization of firms. Based on this premise, this paper explored how the rise of modular production has shaped the pattern of communication and the flow of information inside of firms. We conclude by suggesting several avenues for future research.

One avenue is to explore the organization of firms with non-modular production functions, especially that of multidivisional firms. Alfred Chandler documented the central role multidivisional firms have played in the development of the US economy and spurred a large literature examining their organization (Chandler 1962). A goal of this literature is to understand the firms' choice between *M*, *U*, and matrix forms, between organizing by product, function, or a combination of both (see, for instance, Maskin, Qian, and Xu (2000)). Our paper suggests that this choice is shaped by the firms' production functions. These production functions often have an *overlapping community structure* rather than a non-overlapping one. The R&D decisions for one product, say, have to be closely coordinated with both the manufacturing and marketing decisions for the same product and the R&D decisions for the firm's other products. Even though such production functions are not modular, they fall within the class of general production functions we examined in the extensions in Section 7. The fact that the separability result continues to hold serves as a useful starting point for an exploration of when M, U, and matrix forms are optimal and what determines the choice among them.

A second avenue for future research is to explore the broader impact of modular production on the organization of firms. As we noted in the introduction, Baldwin and Clark (1997) 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 (see Footnote 8). There are many reasons why modular production may affect the boundaries of firms and the structure of industries. We leave their exploration for future research.

A final avenue for future research goes beyond the impact of modular production on organization and asks what explains its rise in the first place. Herbert Simon argued that modularity facilitates adaptation by confining adaptive changes to individual modules within a system (Simon 1962). In line with this intuition, 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.

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] 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.

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

[14] Chandler, Alfred. 1962. Strategy and Structure: Chapters in the History of the Industrial Empire. Cambridge Mass.

[15] Cloteaux, Brian, Michael Drew LaMar, Elizabeth Moseman, and James Shook. 2014. Threshold Digraphs. *National Institute of Standards and Technology.* 119. 227-34.

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

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

[18] Cremer, Jacques, Luis Garicano, and Andrea Prat. 2007. Language and the Theory of the Firm. *Quarterly Journal of Economics*. 122(1): 373–407.

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

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

[21] Eppinger, Steven and Tyson Browning. 2012. Design Structure Matrix Methods and Applications. MIT Press.

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

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

[24] ——— 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.

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

[26] Girvan, Michelle and Mark Newman. 2002. Community Structure in Social and Biological Networks. *Proceedings of the National Academy of Sciences*. 99(12): 7821–26.

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

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

[29] 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.

[30] 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.

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

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

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

[34] Lambert, Nicolas, Giorgio Martini, and Michael ostrovsky. 2018. Quadratic Games. NBER Working Paper Series, No. 24914.

[35] 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.

[36] March, James and Herbert Simon. 1958. Organizations. John Wiley & Sons.

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

[38] Maskin, Eric, Yingyi Qian, and Chenggang Xu. 2000. Incentives, Information, and Organizational Form. *Review of Economic Studies.* 67 (2): 359–78.

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

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

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

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

[43] 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.

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

[45] Rantakari, Heikki. 2008. Governing Adaptation. *Review of Economic Studies*. 75(4): 1257–85.

[46] 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.

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

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

[49] 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.

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

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

[52] 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, $\mathrm{E}_i[Z] \equiv \mathrm{E}\left[Z | \boldsymbol{C}_{(i)}\right]$. Next, a strategy for agent i is a mapping $\tilde{d}_i : \left[-\overline{\theta}, \overline{\theta}\right]^N \to \left[-\overline{D}, \overline{D}\right]$, where $\tilde{d}_i(\theta)$ denotes the decision that agent i makes in state θ. We denote a strategy profile by $\tilde{d} = \times_{i=1}^N \tilde{d}_i$. To ensure equilibrium strategies involve interior decisions, define $\overline{p} = \max_i \sum_j p_{ij}$, and assume that $\overline{D} \geq \frac{\overline{\theta}}{1-\overline{p}}$.

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

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

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

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 $S_i(k)$ as the set of i's pure strategies surviving k rounds of iterated elimination of strictly dominated strategies. Since $d_i(\theta) \in \left[-\overline{D}, \overline{D}\right]$, the first set in the sequence is

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

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,

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

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

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

where

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

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,

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

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

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

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 j_m does 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 decision* d_i^* *puts on its state* θ_i *satisfies* $\omega_{ii}(\boldsymbol{I}_i)a_i = a_i$, *where* $\boldsymbol{I}_i$ *is the ith row of an* $N \times N$ *identity matrix, and it is increasing and supermodular in* $\boldsymbol{C}_i$.

Proof of Corollary 1. This result follows from the proofs of Lemma 1 and Proposition 7 in appendix D. ■

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

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

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

Proof of Lemma 2. Given equilibrium 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 expected revenue can be written as

$$\sum_{i=1}^{N} a_i\mathrm{E}[d_i^*\theta_i] - \sum_{i=1}^{N}\sum_{j=1}^{N} p_{ij}\mathrm{E}\left[d_i^*\left[\mathrm{E}_i\left[d_j^*\right] - d_j^*\right]\right] = \sum_{i=1}^{N} a_i\mathrm{E}[d_i^*\theta_i] - \sum_{i=1}^{N}\sum_{j=1}^{N} p_{ij}\mathrm{E}\left[\mathrm{E}_i[d_i^*]\underbrace{\left[\mathrm{E}_i\left[d_j^*\right] - \mathrm{E}_i\left[d_j^*\right]\right]}_{=0}\right],$$

where the first equality holds by the law of iterated expectations. The result follows because $\mathrm{Cov}(d_i^*, \theta_i) = \mathrm{E}[d_i^* \theta_i] - \mathrm{E}[d_i^*]\,\mathrm{E}[\theta_i]$, and $\mathrm{E}[\theta_i] = 0$ for all i. ■

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

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

for each $i \in \mathcal{N}$, where m_{ij} is a dummy variable that is equal to one if agents i and j 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

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

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 the agents in modules $1, \ldots, \ell$, where $\ell \in \mathcal{M}$, and to none of the agents in the other modules $\mathcal{M} \backslash \{1, \ldots, \ell\}$. Agent 1's expected revenue is then given by*

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

where

$$x_m = \frac{1}{1 + p - (n_m - 1)(p_m - p)} \text{ for } m \in \mathcal{M}.$$

Proof of Lemma 3. 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 $\omega_{11}(\boldsymbol{C}_1(\ell))$ is the $(1,1)$ element of a matrix $\mathbf{Q}^{-1}$.*

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_0. 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.

$$\begin{aligned}
v_0 &= 1 + p_1(n_1-1)v_1 + pn_2v_2 + \cdots + pn_\ell v_\ell \\
v_1 &= p_1v_0 + p_1(n_1-2)v_1 + pn_2v_2 + \cdots + pn_\ell v_\ell \\
&\ \vdots \\
v_\ell &= pv_0 + p(n_1-1)v_1 + pn_2v_2 + \cdots + p_\ell(n_\ell-1)v_\ell.
\end{aligned}$$

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 1; the $k+1$th term, pn_kv_k, is the value of all walks that initially pass to one of the n_k nodes in module k. The right-hand side of the second equation captures the value of all walks from a node $j \neq 1$ in module 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 1; the $k+1$th term is the value of all walks that initially pass to one of the n_k nodes in module k. The remaining equations are interpreted analogously.

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

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & -p_1(n_1-1) & -pn_2 & \cdots & -pn_\ell \\ -p_1 & 1-p_1(n_1-2) & -pn_2 & \cdots & -pn_\ell \\ -p & -p(n_1-1) & 1-p_2(n_2-1) & \cdots & -pn_\ell \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -p & -p(n_1-1) & -pn_2 & \cdots & 1-p_\ell(n_\ell-1) \end{bmatrix} \begin{bmatrix} v_0 \\ v_1 \\ v_2 \\ \vdots \\ v_\ell \end{bmatrix}.$$

Denote this $(\ell+1)\times(\ell+1)$ matrix by $\boldsymbol{Q}$. Then v_0 is the $(1,1)$ element of the inverse matrix $\boldsymbol{Q}^{-1}$, and by the definition of a matrix inverse, $v_0 = \det\tilde{\mathbf{Q}}/\det\mathbf{Q}$, where $\tilde{\mathbf{Q}}$ is the matrix obtained by removing the first row and column of $\boldsymbol{Q}$.

Step 2: *Show that* $\det\mathbf{Q} = \frac{(1+p_1)(1-p_1(n_1-1))}{x_2\cdots x_\ell}\left(1-\lambda\sum_{m=2}^{\ell} n_m x_m\right)$, *where* $\lambda = \frac{p}{1-pn_1x_1}$.

We can write $\mathbf{Q}$ in block-matrix form $\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}$, where

$$\begin{aligned}
\mathbf{A} &= \begin{bmatrix} 1 & -p_1(n_1-1) \\ -p_1 & 1-p_1(n_1-2) \end{bmatrix}, \ \mathbf{B} = \begin{bmatrix} -pn_2 & \cdots & -pn_\ell \\ -pn_2 & \cdots & -pn_\ell \end{bmatrix} \\
\mathbf{C} &= \begin{bmatrix} -p & -p(n_1-1) \\ \vdots & \vdots \\ -p & -p(n_1-1) \end{bmatrix}, \ \mathbf{D} = \begin{bmatrix} 1-p_2(n_2-1) & \cdots & -pn_\ell \\ \vdots & \ddots & \vdots \\ -pn_2 & \cdots & 1-p_\ell(n_\ell-1) \end{bmatrix}.
\end{aligned}$$

By the block matrix determinant formula, $\det \mathbf{Q} = \det(\mathbf{A}) \det\left(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}\right)$. We first calculate $\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$. The expression for $\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$ can be written as the sum of a diagonal matrix and a rank-one matrix:

$$\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B} = \mathbf{X}^{-1} - \lambda \mathbf{u}\mathbf{v}^T,$$

where $\lambda = \frac{p}{1-pn_1x_1}$, and

$$\mathbf{X}^{-1} = \begin{bmatrix} x_2^{-1} & \cdots & 0 \\ \vdots & \vdots & \vdots \\ 0 & \cdots & x_\ell^{-1} \end{bmatrix}, \mathbf{u} = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}, \mathbf{v} = \begin{bmatrix} n_2 \\ \vdots \\ n_\ell \end{bmatrix}.$$

By the matrix determinant lemma,

$$\det\left(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}\right) = \left(1 - \lambda \mathbf{v}^T \mathbf{X}\mathbf{u}\right) \det \mathbf{X}^{-1} = \frac{1 - \lambda \sum_{m=2}^{\ell} n_m x_m}{x_2 \cdots x_\ell}.$$

We therefore have that

$$\det \mathbf{Q} = \frac{(1+p_1)(1-p_1(n_1-1))}{x_2 \cdots x_\ell} \left(1 - \lambda \sum_{m=2}^{\ell} n_m x_m\right).$$

Step 3: *Show that* $\det \tilde{\mathbf{Q}} = \frac{1-p_1(n_1-2)}{x_2 \cdots x_\ell}\left(1 - \tilde{\lambda}\sum_{m=2}^{\ell} n_m x_m\right)$, *where* $\tilde{\lambda} = \frac{p}{1-p(n_1-1)\tilde{x}_1}$ *and* $\tilde{x}_1 = \frac{1}{1-p_1(n_1-2)+p(n_1-1)}$.

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

$$\tilde{\mathbf{D}} - \tilde{\mathbf{C}}\tilde{\mathbf{A}}^{-1}\tilde{\mathbf{B}} = \mathbf{X}^{-1} - \tilde{\lambda}\mathbf{u}\mathbf{v}^T,$$

where $\mathbf{X}^{-1}$, $\mathbf{u}$, and $\mathbf{v}$ are the same as in step 2, $\tilde{\lambda} = \frac{p}{1-p(n_1-1)\tilde{x}_1}$, and $\tilde{x}_1 = \frac{1}{1-p_1(n_1-2)+p(n_1-1)}$. As above, the matrix determinant lemma,

$$\det\left(\tilde{\mathbf{D}} - \tilde{\mathbf{C}}\tilde{\mathbf{A}}^{-1}\tilde{\mathbf{B}}\right) = \left(1 - \tilde{\lambda}\mathbf{v}^T\mathbf{X}\mathbf{u}\right)\det\mathbf{X}^{-1} = \frac{1 - \tilde{\lambda}\sum_{m=2}^{\ell} n_m x_m}{x_2 \cdots x_\ell},$$

then gives us the result.

Step 4. *Show that* $\omega_{11}(\mathbf{C}_1(\ell)) = \frac{1+(p_1-p)x_1}{1+p_1} + \frac{p x_1^2}{1-p\sum_{m=1}^{\ell} n_m x_m}$.

By the preceding three steps,

$$\begin{aligned}\omega_{11}(\mathbf{C}_1(\ell)) &= v_0 = \frac{\det\tilde{\mathbf{Q}}}{\det\mathbf{Q}} = \frac{(1-p_1(n_1-2))\left(1-\tilde{\lambda}\sum_{j=2}^{\ell} n_j x_j\right)}{(1+p_1)(1-p_1(n_1-1))\left(1-\lambda\sum_{j=2}^{\ell} n_j x_j\right)} \\ &= \frac{1+(p_1-p)x_1}{1+p_1} + \frac{p x_1^2}{1-p\sum_{m=1}^{\ell} n_m x_m}.\end{aligned}$$

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. *There exist thresholds $\lambda_i \geq 0$ such that it is optimal for agent $i \in \mathcal{N}$ to tell his state to a different agent $j \in \mathcal{N}$ if and only if:*

(*i.*) *agent j belongs to the same module $m(j) = m(i)$ or*

(*ii.*) *agent j belongs to a different module $m(j) \neq m(i)$ with cohesion $x_{m(j)} \geq \lambda_i$.*

The threshold λ_i is increasing in γ, decreasing in $a_i^2\sigma_i^2$, p_m, and n_m for all $m \in \mathcal{M}$, and independent of $a_k^2\sigma_k^2$ for any $k \in \mathcal{N}\setminus\{i\}$.

Proof of Proposition 2. Let $\mathcal{K}$ be an arbitrary set of modules with $m(i) \in \mathcal{K}$, define $S(\mathcal{K}) = \sum_{m:m\in\mathcal{K}} n_m x_m$, and denote by $\mathbf{C}_i(\mathcal{K})$ the row of the communication matrix in which agent i tells θ_i to agent j if and only if $m(j) \in \mathcal{K}$. By Lemma 3, $R_i(\mathbf{C}_i(\mathcal{K}))$ can be written as an increasing and convex function of $S(\mathcal{K})$, $h(S(\mathcal{K}))$.

Now suppose that it is optimal to inform all modules in $\mathcal{K}$. Then it must be the case that for all m with $m \in \mathcal{K}$

$$\gamma \leq \frac{h(S(\mathcal{K})) - h(S(\mathcal{K}\backslash m))}{n_m} = x_m \frac{h(S(\mathcal{K})) - h(S(\mathcal{K}\backslash m))}{S(\mathcal{K}) - S(\mathcal{K}\backslash m)} < x_m h'(S(\mathcal{K})),$$

where the last inequality holds because h is convex. Suppose further that it is not optimal to also inform some module $m' \notin \mathcal{K}$. Then it must be the case that

$$\gamma > \frac{h(S(\mathcal{K}\cup\{m'\})) - h(S(\mathcal{K}))}{n_{m'}} = x_{m'} \frac{h(S(\mathcal{K}\cup\{m'\})) - h(S(\mathcal{K}))}{S(\mathcal{K}\cup\{m'\}) - S(\mathcal{K})} > x_{m'} h'(S(\mathcal{K})).$$

These two inequalities imply that $x_m > x_{m'}$ for all modules m that are optimally told about θ_i and modules m' that are optimally not told about θ_i. In other words, there is some threshold λ_i such that agent i tells θ_i to agent j in module $m(j) \neq m(i)$ if and only if $x_{m(j)} \geq \lambda_i$.

The threshold λ_i is given by the value at which the expected revenue of informing all modules m with $x_m \geq \lambda_i$ minus the cost of informing those modules is highest. If profits are negative from informing any subset of modules, then $\lambda_i > x_m$ for all $m \in \{2, \ldots, M\}$. Recall that the expected revenue generated by agent i when modules $\{1, \ldots, \ell\}$ know state θ_i is

$$a_i^2 \sigma_i^2 \left(\frac{1 + \left(p_{m(i)} - p\right) x_{m(i)}}{1 + p_{m(i)}} + \frac{p x_{m(i)}^2}{1 - p \sum_{m=1}^{\ell} n_m x_m} \right),$$

where $x_m = \frac{1}{1+p-(n_m-1)(p_m-p)}$. The first term in brackets is increasing in $n_{m(i)}$ and $p_{m(i)}$, and the second term is increasing in $x_{m(i)}$, which is increasing in $p_{m(i)}$ and $n_{m(i)}$, since $p_m \geq p$ for all m by assumption. Thus, the expected revenue from informing modules $\{1, \ldots, \ell\}$ is increasing in $a_{m(i)}, \sigma_{m(i)}^2, n_{m(i)}$, and $p_{m(i)}$. Similarly, the expected per-node incremental revenue from additionally informing modules $\{\ell+1, \ldots, \ell+k\}$ given modules $\{1, \ldots, \ell\}$ are already informed is increasing in $a_{m(i)}, \sigma_{m(i)}^2, n_{m(i)}$, and $p_{m(i)}$. Following an increase in $a_{m(i)}, \sigma_{m(i)}^2, n_{m(i)}$, and $p_{m(i)}$, since the incremental revenue from additionally informing modules is increasing in $a_{m(i)}, \sigma_{m(i)}^2, n_{m(i)}$, and $p_{m(i)}$, then the profit from informing all modules with $x_m \geq \lambda_i$ continues to be higher than the profit from informing a subset of those modules. ■

COROLLARY 2. *Optimal communication networks have the threshold property.*

Proof of Corollary 2. This result follows from Proposition 2 with $s_i = \lambda_i$ and $r_j = x_{m(j)}$ for all i, j. ■

LEMMA 4. *A communication network has the threshold property if and only if it contains no two-switches or directed three-cycles.*

Proof of Lemma 4. To establish this result, we will use Theorem 1 of Cloteaux et al. (2014), which establishes a forbidden subgraph characterization of threshold directed graphs. To do so, we first introduce two definitions.

Let $\mathbf{C}$ be a communication network, and say that $\mathbf{\Psi}$ is *induced by* $\mathbf{C}$ if $\mathbf{\Psi}$ is a subgraph of $\mathbf{C}$ consisting of $\mathbf{M}$ nodes $\{i_1, \ldots, i_M\}$ with $m\left(i_\ell\right) = \ell$ and m, m' entry $\psi_{mm'} = c_{i_m i_{m'}}$ if $m \neq m'$ and $\psi_{mm} = 0$. The network $\mathbf{\Psi}$ is a *threshold directed graph* if for all m, m', m'' distinct, if $\sum_\ell \psi_{m\ell} \geq \sum_\ell \psi_{m'\ell}$ (and $\sum_\ell \psi_{\ell m} \geq \sum_\ell \psi_{\ell m'}$ if $\sum_\ell \psi_{m\ell} = \sum_\ell \psi_{m'\ell}$), then $\psi_{m'm''} = 1$ implies $\psi_{mm''} = 1$.

We begin with an alternate characterization of threshold directed graphs that parallels our threshold condition.

Step 1: *Show* $\mathbf{\Psi}$ *is a threshold directed graph if and only if there exists two sequences of nonnegative real numbers* $\{\tilde{s}_1, \ldots, \tilde{s}_M\}$ *and* $\{\tilde{r}_1, \ldots, \tilde{r}_M\}$ *such that* $\psi_{mm'} = 1$ *if and only if* $\tilde{s}_m \leq \tilde{r}_{m'}$.

Suppose there exists $\{\tilde{s}_1, \ldots, \tilde{s}_M\}$ and $\{\tilde{r}_1, \ldots, \tilde{r}_M\}$ such that $\psi_{mm'} = 1$ if and only if $\tilde{s}_m \leq \tilde{r}_{m'}$.

Arrange the nodes so that $\tilde{s}_1 \leq \tilde{s}_2 \leq \cdots \leq \tilde{s}_M$, and $\tilde{s}_m = \tilde{s}_{m+1}$ implies $\tilde{r}_m \geq \tilde{r}_{m+1}$, and define $\overline{\beta}_m = \max_{\ell \neq m} \{\ell | \, \tilde{s}_\ell \leq \tilde{r}_m\}$. Let

$$\beta_m = \sum_{\ell=1}^{M} \psi_{\ell m} = \begin{cases} \overline{\beta}_m & \text{if } \overline{\beta}_m \leq m \\ \overline{\beta}_m - 1 & \text{if } \overline{\beta}_m > m. \end{cases}$$

Then $\psi_{mm'} = 1$ if and only if $\tilde{s}_m \leq \tilde{r}_{m'}$ if and only if $m \leq \beta_{m'}$ if $m < m'$ and $m \leq \beta_{m'} + 1$ if $m > m'$. We therefore have that

$$\psi_{mm'} = \begin{cases} 1 & \text{if } m < m' \text{ and } m \leq \beta_{m'} \\ 1 & \text{if } m > m' \text{ and } m \leq \beta_{m'} + 1 \\ 0 & \text{otherwise,} \end{cases}$$

and therefore by Corollary 3 of Cloteaux et al. (2014), Ψ is a threshold directed graph.

Conversely, suppose $\mathbf{\Psi}$ is a threshold directed graph. If we arrange the nodes so that $\sum_\ell \psi_{m\ell} \geq \sum_\ell \psi_{m+1,\ell}$ (and $\sum_\ell \psi_{\ell m} \geq \sum_\ell \psi_{\ell,m+1}$ if $\sum_\ell \psi_{m\ell} = \sum_\ell \psi_{m+1,\ell}$) and let $\beta_m = \sum_\ell \psi_{\ell m}$, then by Corollary 3 of Cloteaux et al. (2014),

$$\psi_{mm'} = \begin{cases} 1 & \text{if } m < m' \text{ and } m \leq \beta_{m'} \\ 1 & \text{if } m > m' \text{ and } m \leq \beta_{m'} + 1 \\ 0 & \text{otherwise.} \end{cases}$$

Let $\tilde{s}_m = m$ and $\tilde{r}_m = \min_\ell \{\ell | \, \psi_{\ell m} = 1\}$. Then $\psi_{mm'} = 1$ if $\tilde{s}_m \leq \tilde{r}_{m'}$ and $\psi_{mm'} = 0$ otherwise, establishing the claim.

Step 2: *Show that if* **C** *has the threshold property, for all* $\mathbf{\Psi}$ *induced by* **C**, $\mathbf{\Psi}$ *is a threshold directed graph.*

This step follows from step 1 by setting $\tilde{s}_m = s_{i_m}$ and $\tilde{r}_m = r_{i_m}$ for all m.

Step 3: *Argue that if* **C** *has the threshold property, it has no two-switches or directed three-cycles.*

By Theorem 1 of Cloteaux et al. (2014), $\mathbf{\Psi}$ is a threshold directed graph if and only if it has no two-switches or directed three-cycles. Suppose **C** has a two-switch or a directed three-cycle. Then consider $\mathbf{\Psi}$ induced by **C** containing the two-switch or the directed three-cycle. Then $\mathbf{\Psi}$ is not a threshold directed graph, so **C** does not have the threshold property. ■

COROLLARY 3. *A communication network is not optimal if it is a tree or matrix.*

Proof of Corollary 3. Suppose **C** is a tree. Let i denote an agent in the module in level 1, j and k be agents in two different modules in level 2, and ℓ be an agent in level 3 whose unique predecessor module is $m(j)$. This is without loss of generality, since levels 2 and 3 have at least two modules with a unique predecessor module. Then agent i tells his state to agent k but not to agent ℓ, and agent j tells his state to agent ℓ but not to agent k. The communication network $\mathbf{C}$ therefore contains a two-switch and by Corollary 2 is not optimal.

Next, suppose $\mathbf{C}$ is a matrix**.** Consider two agents i and k in different modules but in the same horizontal team. Consider a different horizontal team containing two agents j and ℓ in two different modules. Since any two modules in the same horizontal team are in different vertical teams, then of the four modules containing i, j, k, and ℓ, either these modules are partitioned into two vertical teams, or they are partitioned into more than two vertical teams.

Suppose they are partitioned into two vertical teams. Suppose agent i's module and agent j's module are in one vertical team, and agent k's module and agent ℓ's module are in another vertical team. Then i tells his state to k but not ℓ, and j tells his state to ℓ but not k, so $\mathbf{C}$ contains a two-switch and is therefore not optimal. Suppose instead i, j, k, and ℓ are partitioned into at least three vertical teams. Then at least one module from each horizontal team, the modules containing i and j, say, must be in a vertical team without any of the other four modules containing i, j, k, or ℓ. Then i tells his state to k but not ℓ, and j tells his state to ℓ but not k, so $\mathbf{C}$ contains a two-switch and is therefore not optimal. ■

PROPOSITION 3. *Suppose that for any agents $i, j \in \mathcal{N}$, $\lambda_i \leq \lambda_j$ if and only if $x_{m(i)} \geq x_{m(j)}$. Any optimal communication network then has a core-periphery structure in which the agents who belong to the most cohesive modules form the core.*

Proof of Proposition 3. We construct a partition of the set of agents that has a core-periphery structure as described. Label the modules by their cohesion, with the most cohesive labeled 1 and the least cohesive labeled 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 j in module k tells his state to all agents in module $k-1$. The ordering of modules by cohesion is unique up to modules with the same cohesion. We deal with the case where no such k exists below.

We first show that agents in modules $1, \ldots, k-1$ are in the core. By Proposition 2, agent j tells his state to all agents in modules $1, \ldots, k-1$ and, since $\lambda_j \geq \lambda_i$ for all agents i in modules weakly more cohesive than j's module,then all agents in modules $1, \ldots, k-1$ tell their state to all others in modules $1, \ldots, k-1$. Thus we say modules $1, \ldots, k-1$ are in the core.

Next we examine whether agents in module k are in the core. If all agents in module $k-1$ tell their state to agents in module k, since $\lambda_i \leq \lambda_j$ if and only if $x_{m(i)} \geq x_{m(j)}$, then all agents in modules $1, \ldots, k-1$ tell their state to all agents in module k. Then any agents in module k that tell their state to agents in modules $1, \ldots, k-1$ are in the core. By Proposition 2, any agents in module k that do not tell their state to agents in module $k-1$ also do not inform any agents in modules $k+1, \ldots, M$, and we say they are in the periphery. If instead some agents in module $k-1$ do not tell their state to agents in module k then, by Proposition 2, those agents in module $k-1$ do not tell their state to any agents in modules $k, \ldots, M$. Since $\lambda_i \leq \lambda_j$ if and only if $x_{m(i)} \geq x_{m(j)}$, then agents in module k also do not tell their state to any agents in modules $k+1, \ldots, M$ and we say they are in the periphery.

We next show that agents in modules $k+1, \ldots, M$ do not tell their states to agents in modules $k, \ldots, M$ outside their own module, and we say they are in the periphery. By definition of module k, there is no agent in module $k+1$ who tells his state to agents in module k. Then by Proposition 2 and since $\lambda_i \leq \lambda_j$ if and only if $x_{m(i)} \geq x_{m(j)}$, no agent in module $k+1, \ldots, M$ tells his state to any agents in modules $k, \ldots, M$, aside from those agents in his own module.

If there is no $k \in \{2, \ldots, M\}$ such that some agent j in module k tells his state to all agents in module $k-1$, then no agent in module 2 tells his state to any agent in module 1. By Proposition 2 and since $\lambda_i \leq \lambda_j$ if and only if $x_{m(i)} \geq x_{m(j)}$, then no agent in modules $2, \ldots, M$ tells their state to any agent outside their own module. We say agents in module 1 are in the core, and all other agents are in the periphery. ■

For the next lemma, we introduce a definition.

DEFINITION. *Let* $\mathbf{\Psi}$ *be a directed graph consisting of* M *nodes. We say that* $\mathbf{\Psi}$ *has a "generalized core-periphery structure" if its nodes can be partitioned into a core (denoted* $\mathcal{M}^C$*), a periphery (denoted* $\mathcal{M}^P$*), and a suburban periphery (denoted* $\mathcal{M}^{SP}$*), at least two of which are non-empty, satisfying the following properties:*

(i.) the core is a clique: $\psi_{mm'} = 1$ *for all* $m, m' \in \mathcal{M}^C$*,*

(ii.) the periphery is an independent set: $\psi_{mm'} = 0$ *for all* $m, m' \in \mathcal{M}^P$*, and*

(iii.) for each $m \in \mathcal{M}^{SP}$*, there exists some* $m' \in \mathcal{M}^C$ *(if non-empty) and* $m'' \in \mathcal{M}^P$ *(if non-empty) such that either (a)* $\psi_{m'm} = 0$ *and* $\psi_{m''m} = 1$ *or (b)* $\psi_{mm'} = 0$ *and* $\psi_{mm''} = 1$.

LEMMA 5. *Let* $\mathbf{C}$ *be an optimal communication network, and suppose* $\mathbf{\Psi}$ *is induced by* $\mathbf{C}$*. Then* $\mathbf{\Psi}$ *has a generalized core-periphery structure.*

Proof of Lemma 5. For the definition of what it means for $\mathbf{\Psi}$ to be induced by $\mathbf{C}$, see the proof of Lemma 4. Arrange the nodes $\{i_1, \ldots, i_M\}$ in Ψ so that $\lambda_1 \leq \cdots \leq \lambda_M$, and consider a trivial partition of $\{i_1, \ldots, i_M\}$ in which $m(i_\ell) \in \mathcal{M}^{SP}$ for all ℓ. There are two cases two consider. In the first case, suppose for all m, $\psi_{i_M i_m} = 0$. Consider an alternative partition in which $m(i_\ell) \in \mathcal{M}^{SP}$ for all $\ell < M$, and $m(i_M) \in \mathcal{M}^P$. In the second case, suppose $\psi_{i_M i_m} = 1$ for some $m < M$. Then take $\overline{m} = \text{argmax}_{m<M} x_m$. By construction, $\psi_{i_M i_{\overline{m}}} = 1$ and therefore $\psi_{i_m i_{\overline{m}}} = 1$ for all $m \neq \overline{m}$.

Consider an alternative partition in which $m(i_\ell) \in \mathcal{M}^{SP}$ for all $\ell \neq \overline{m}$, and $\overline{m} \in \mathcal{M}^C$. In both cases, therefore, $\mathbf{\Psi}$ has a partition into a core, periphery, and suburban periphery, at least two of which are non-empty, and therefore $\mathbf{\Psi}$ has a generalized core-periphery structure. ■

PROPOSITION 4. *If production network $\boldsymbol{P}$ is weakly modular, it does not admit any fragmented communication networks.*

Proof of Proposition 4. We will show that if $\mathbf{P}$ is weakly modular, then the per-node return to informing another module about θ_i is always increasing in the number of modules whose agents know θ_i. This result implies that it is either optimal for i to tell his state only to agents in module $m(i)$ or to tell his state to all agents.

To make this argument, consider agent 1, and order the remaining modules in decreasing order of their cohesion: $x_2 \geq x_3 \geq \cdots \geq x_M$. Suppose modules $\{1, \ldots, \ell-1\}$ are told about θ_1, and consider the per-node returns to informing module ℓ. If we define $S_\ell \equiv \sum_{m=1}^{\ell} n_m x_m$, then by Lemma 3, we can write the expected revenues $R_1(\mathbf{C}_1(\ell)) = a_1^2\sigma_1^2\left(\frac{1+(p_1-p)x_1}{1+p_1} + \frac{px_1^2}{1-pS_\ell}\right)$, so the per-node returns to informing module ℓ are

$$\frac{R_1(\mathbf{C}_1(\ell)) - R_1(\mathbf{C}_1(\ell-1))}{n_\ell} = x_\ell \frac{a_1^2\sigma_1^2 p^2 x_1^2}{(1-pS_\ell)(1-pS_{\ell-1})} \equiv x_\ell b_\ell.$$

Next, note that $x_\ell b_\ell < x_{\ell+1} b_{\ell+1}$ if and only if

$$x_\ell - x_{\ell+1} < \frac{n_\ell x_\ell + n_{\ell+1} x_{\ell+1}}{1-pS_{\ell+1}} p x_{\ell+1}.$$

Since $x_m \geq \frac{1}{1+p}$ for all m, the left-hand side of this inequality is less than $x_\ell - \frac{1}{1+p}$. And since $n_m \geq 1$ for all m, the right-hand side is greater than $\frac{2}{1+p}\frac{p}{1+p}$. We therefore have that a sufficient condition for $x_\ell b_\ell < x_{\ell+1} b_{\ell+1}$ for all ℓ is that

$$x_\ell - \frac{1}{1+p} < \frac{2}{1+p}\frac{p}{1+p}.$$

This inequality is satisfied for all ℓ if

$$p_\ell - p < \frac{2p(1+p)}{(1+3p)(n_\ell - 1)},$$

that is, if $\mathbf{P}$ is weakly modular. ■

PROPOSITION 5. *If production network $\boldsymbol{P}$ is weakly modular, and the modules differ in their cohesions, there exists a $\overline{\mu} \in (0,1)$ such that $\boldsymbol{P}(\mu)$ admits a fragmented communication network for all $\mu \in [\overline{\mu}, 1)$.*

Proof of Proposition 5. Recall from the proof of Proposition 4 that if we order the modules in decreasing order of their cohesion, we can write the per-node returns to informing module ℓ as

$$\frac{R_1\left(\mathbf{C}_1\left(\ell\right)\right)-R_1\left(\mathbf{C}_1\left(\ell-1\right)\right)}{n_\ell}=x_\ell\frac{a_1^2\sigma_1^2p^2x_1^2}{\left(1-pS_\ell\right)\left(1-pS_{\ell-1}\right)}\equiv x_\ell b_\ell.$$

Suppose $\mathbf{P}$ is weakly modular, and denote by $x_\ell\left(\mu\right)$, $b_\ell\left(\mu\right)$, and $S_\ell\left(\mu\right)$ the objects defined above that correspond to the weight-neutral increase in modularity of size μ of $\mathbf{P}$. We will show that there is a $\overline{\mu}$ such that if $\mu\geq\overline{\mu}$, $x_\ell\left(\mu\right)b_\ell\left(\mu\right)>x_{\ell+1}\left(\mu\right)b_{\ell+1}\left(\mu\right)$, that is, the per-node returns to informing another module about a state θ_i is always decreasing in the number of modules whose agents know θ_i. To this end, note that $x_{\ell+1}\left(\mu\right)b_{\ell+1}\left(\mu\right)-x_\ell\left(\mu\right)b_\ell\left(\mu\right)$ is equal to

$$\frac{a_1^2\sigma_1^2p\left(\mu\right)^2x_1\left(\mu\right)^2}{\left(1-p\left(\mu\right)S_{\ell+1}\left(\mu\right)\right)\left(1-p\left(\mu\right)S_\ell\left(\mu\right)\right)}\underbrace{\left[\frac{n_\ell x_\ell\left(\mu\right)+n_{\ell+1}x_{\ell+1}\left(\mu\right)}{1-p\left(\mu\right)S_{\ell-1}\left(\mu\right)}p\left(\mu\right)x_\ell\left(\mu\right)-\left(x_\ell\left(\mu\right)-x_{\ell+1}\left(\mu\right)\right)\right]}_{\equiv H_\ell(\mu)}.$$

By the proof of Proposition 4, since $\mathbf{P}$ is weakly modular, $H_\ell\left(0\right)>0$ for all ℓ. Moreover, $H_\ell\left(1\right)\leq 0$ for all ℓ. To see why, note that $x_\ell>x_{\ell+1}$ and $\frac{d}{d\mu}\left(x_\ell\left(\mu\right)-x_{\ell+1}\left(\mu\right)\right)=pN\left(x_\ell\left(\mu\right)^2-x_{\ell+1}\left(\mu\right)^2\right)>0$ implies $x_\ell\left(1\right)>x_{\ell+1}\left(1\right)$. We therefore have that $H_\ell\left(1\right)=-\left(x_\ell\left(1\right)-x_{\ell+1}\left(1\right)\right)<0$. Let $\overline{\mu}=\sup_{\mu\in(0,1),\ell}\left\{H_\ell\left(\mu\right)\right\}$. Since $H_\ell\left(\mu\right)$ is continuous in μ, we have that $\overline{\mu}<1$, and for all $\mu>\overline{\mu}$, $x_\ell\left(\mu\right)b_\ell\left(\mu\right)>x_{\ell+1}\left(\mu\right)b_{\ell+1}\left(\mu\right)$ for all ℓ, so $\mathbf{P}\left(\mu\right)$ admits a fragmented communication network. ■

Appendix B: Costly Within-Module Communication

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 $\tilde{R}_1\left(\boldsymbol{C}_1\left(\ell\right)\right)$ agent 1's expected revenue if he tells his state to agents in modules $2,\ldots,\ell$ but not to agents in his own module 1. The following lemma derives an expression for $\tilde{R}_1\left(\boldsymbol{C}_1\left(\ell\right)\right)$ in terms of the model primitives.

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

$$\tilde{R}_1\left(\boldsymbol{C}_1\left(\ell\right)\right)=a_1^2\sigma_1^2\left(1+\frac{p^2\left(\sum_{m=1}^{\ell}n_mx_m-n_1x_1\right)}{1-\left(p+p^2\right)\sum_{m=2}^{\ell}n_mx_m}\right),$$

where

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

Proof of Lemma B1. The proof of this lemma parallels the proof of Lemma 3, but for the case where agents in module 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 v_0. 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 as a system of equations

$$\begin{aligned}
v_0 &= 1 + pn_2v_2 + \cdots + pn_\ell v_\ell \\
v_2 &= pv_0 + p_2(n_2-1)v_2 + pn_3v_3 + \cdots + pn_\ell v_\ell \\
&\vdots \\
v_\ell &= pv_0 + pn_2v_2 + pn_3v_3 + \cdots + p_\ell(n_\ell-1)v_\ell.
\end{aligned}$$

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

$$\begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & -pn_2 & \cdots & pn_\ell \\ -p & 1-p_2(n_2-1) & \cdots & pn_\ell \\ \vdots & \vdots & \ddots & \vdots \\ -p & -pn_2 & \cdots & 1-p_\ell(n_\ell-1) \end{bmatrix} \begin{bmatrix} v_0 \\ v_2 \\ \vdots \\ v_\ell \end{bmatrix}.$$

We denote this $\ell \times \ell$ matrix by $\boldsymbol{Q}$. Then v_0 is the $(1,1)$ element of the inverse matrix $\boldsymbol{Q}^{-1}$, and by the definition of a matrix inverse, $v_0 = \det \tilde{\boldsymbol{Q}} / \det \boldsymbol{Q}$, where $\tilde{\boldsymbol{Q}}$ is the matrix obtained by removing the first row and column of $\boldsymbol{Q}$.

Next, we derive $\det \boldsymbol{Q}$ 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 $\boldsymbol{Q}$ such that $\boldsymbol{A} = 1$, and $\boldsymbol{B}$, $\boldsymbol{C}$, and $\boldsymbol{D}$ are defined accordingly. Then

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

Define $x_m^{-1} = 1 - p_m(n_m - 1) + pn_m$ and $z = p(1+p)$. 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 $\det\left(\boldsymbol{D}-\boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B}\right) = \frac{1}{x_2 \cdots x_\ell}\left(1 - z\sum_{m=2}^{\ell} n_m x_m\right)$ for all $\ell \geq 2$. Since $\det \boldsymbol{A} = 1$, we have

$$\det \boldsymbol{Q} = \frac{1}{x_2 \cdots x_\ell}\left(1 - p\left(1+p\right)\sum_{m=2}^{\ell} n_m x_m\right) \text{ for all } \ell \geq 2.$$

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

$$\det \tilde{\boldsymbol{Q}} = \frac{1}{x_2 \cdots x_\ell}\left(1 - p\sum_{m=2}^{\ell} n_m x_m\right) \text{ for all } \ell \geq 2.$$

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

$$\tilde{\omega}_{11}\left(\boldsymbol{C}_1\left(\ell\right)\right) = \frac{\det \tilde{\boldsymbol{Q}}}{\det \boldsymbol{Q}} = 1 + \frac{p^2\sum_{m=2}^{\ell} n_m x_m}{1-\left(p+p^2\right)\sum_{m=2}^{\ell} n_m x_m},$$

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.*

Proof of Proposition B1. Let $\mathcal{K}$ be an arbitrary set of modules with $m\left(i\right) \notin \mathcal{K}$, define $S\left(\mathcal{K}\right) = \sum_{m:m\in\mathcal{K}} n_m x_m$, and denote by $\mathbf{C}_i\left(\mathcal{K}\right)$ the row of the communication matrix in which agent i tells θ_i to agent j if and only if $m\left(j\right) \in \mathcal{K}$. By Lemma B1, $\tilde{R}_i\left(\mathbf{C}_i\left(\mathcal{K}\right)\right)$ can be written as an increasing and convex function of $S\left(\mathcal{K}\right)$. The proof of the threshold result then replicates the proof of Proposition 2. The comparative statics on $\overline{\lambda}_i$ follow from Proposition 7. ■

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

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

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. *A 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 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 $\mathrm{E}^{\tilde{\boldsymbol{C}}}\left[\omega_{ii}\left(\tilde{\boldsymbol{C}}_i\right)\middle|\,\boldsymbol{C}\right] = (1-q) + q\omega_{ii}\left(\boldsymbol{C}_i\right)$. 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 $\omega_{ii}\left(\tilde{\boldsymbol{C}}_i\right) = \omega_{ii}\left(\boldsymbol{C}_i\right)$. The principal's objective is therefore separable across agents i, and her objective for agent i is to

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

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

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

which establishes the result. ■

Appendix D: Heterogeneous Coupling

Each node i belongs to a module $m(i) \in \mathcal{M}$, and each module m belongs to a cluster $k(m) \in \mathcal{K}$, where $\mathcal{K} = \{1, \ldots, K\}$. As in the main model, the need for coordination between two decisions depends on whether they are in the same module. In contrast to the main model, the need for coordination between two decisions in different modules depends on whether they are in the same cluster. That is, $p_{ij} = p_m$ if $m(i) = m(j)$, $p_{ij} = p^k$ if $m(i) \neq m(j)$ but $k(m(i)) = k(m(j))$, and $p_{ij} = p$ otherwise. Note that throughout this extension, we will use subscripts to denote module-level characteristics and superscripts to denote cluster-level characteristics.

The proof of this proposition proceeds by establishing two results, which parallel Lemma 3 and Proposition 2. First, Lemma D1 derives a closed-form expression for the expected revenues that result when node 1 informs an arbitrary set of modules in an arbitrary set of clusters. Then, Proposition 6 uses the convexity of the resulting function in each of its arguments to show that it implies a cluster-level threshold property for optimal communication networks.

LEMMA D1. *Suppose agent 1 tells his state to all agents in an arbitrary set of modules $\mathcal{M}^*$ that includes module $m(1)$. Agent 1's expected revenue is then given by*

$$R_1(\mathbf{C}_1) = a_1^2\sigma_1^2\left(\frac{1+\left(p_1-p^1\right)x_1}{1+p_1} + \frac{\left(p^1-p\right)x_1^2}{1-\left(p^1-p\right)S^1} + \left(\frac{1}{1-\left(p^1-p\right)S^1}\right)^2 \frac{px_1^2}{1-p\sum_{k=1}^{K}\frac{S^k}{1-\left(p^k-p\right)S^k}}\right),$$

where

$$S^k = \sum_{m\in\mathcal{M}^*,k=k(m)} n_m x_m \text{ and } x_m = \frac{1}{1+p-(n_m-1)(p_m-p)}.$$

Proof of Lemma D1. Suppose agent 1 is in module $m=1$ in cluster $k=1$, and suppose $\mathcal{M}^*$ contains k_1 modules in cluster $k=1$, k_2-k_1 modules in cluster $k=2$, and k_K-k_{K-1} modules in cluster $k=K$. Number these modules so that modules $m=1$ to $m=k_1$ are in cluster 1, modules $m=k_1+1$ to $m=k_2$ are in cluster 2, and modules $m=k_{K-1}+1$ to $m=k_K$ are in cluster K.

By the same argument as in step 1 of the proof of Lemma 3, if we let v_0 represent the sum of the value 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 and v_ℓ the sum of the values of all walks from a node in module ℓ to node 1 on this same subgraph, then

$$\begin{bmatrix}1\\0\\\vdots\\0\end{bmatrix} = \mathbf{Q}\begin{bmatrix}v_0\\v_1\\\vdots\\v_{n_{k_K}}\end{bmatrix},$$

where

$$\boldsymbol{Q} = \begin{array}{l}\quad\ \overbrace{\qquad\qquad\qquad}^{\text{module }1}\ \ \overbrace{\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad}^{\text{cluster }1}\ \overbrace{\qquad\qquad\qquad\qquad\qquad}^{\text{cluster }2}\qquad\quad \overbrace{\qquad\qquad\qquad\qquad\qquad}^{\text{cluster }K}\\ \begin{bmatrix}
1 & -p_1(n_1-1) & -p^1n_2 & \cdots & -p^1n_{k_1} & -pn_{k_1+1} & \cdots & -pn_{k_2} & \cdots & -pn_{k_{K-1}+1} & \cdots & -pn_{k_K}\\
-p_1 & 1-p_1(n_1-2) & -p^1n_2 & \cdots & -p^1n_{k_1} & -pn_{k_1+1} & \cdots & -pn_{k_2} & \cdots & -pn_{k_{K-1}+1} & \cdots & -pn_{k_K}\\
-p^1 & -p^1(n_1-1) & x_2^{-1}-p^1n_2 & \cdots & -p^1n_{k_1} & -pn_{k_1+1} & \cdots & -pn_{k_2} & \cdots & -pn_{k_{K-1}+1} & \cdots & -pn_{k_K}\\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots\\
-p^1 & -p^1(n_1-1) & -p^1n_2 & \cdots & x_{k_1}^{-1}-p^1n_{k_1} & -pn_{k_1+1} & \cdots & -pn_{k_2} & \cdots & -pn_{k_{K-1}+1} & \cdots & -pn_{k_K}\\
-p & -p(n_1-1) & -pn_2 & \cdots & -pn_{k_1} & x_{k_1+1}^{-1}-p^2n_{k_1+1} & \cdots & -p^2n_{k_2} & \cdots & -pn_{k_{K-1}+1} & \cdots & -pn_{k_K}\\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots\\
-p & -p(n_1-1) & -pn_2 & \cdots & -pn_{k_1} & -p^2n_{k_1+1} & \cdots & x_{k_2}^{-1}-p^2n_{k_2} & \cdots & -pn_{k_{K-1}+1} & \cdots & -pn_{k_K}\\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots\\
-p & -p(n_1-1) & -pn_2 & \cdots & -pn_{k_1} & -pn_{k_1+1} & \cdots & -pn_{k_2} & \cdots & x_{k_{K-1}+1}^{-1}-p^Kn_{k_{K-1}+1} & \cdots & -p^Kn_{k_K}\\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots\\
-p & -p(n_1-1) & -pn_2 & \cdots & -pn_{k_1} & -pn_{k_1+1} & \cdots & -pn_{k_2} & \cdots & -p^Kn_{k_{K-1}+1} & \cdots & x_{k_K}^{-1}-p^Kn_{k_K}
\end{bmatrix}\end{array}$$

If we let $\widetilde{\mathbf{Q}}$ be the matrix obtained by removing the first row and column of $\mathbf{Q}$, then $v_0 = \det \widetilde{\mathbf{Q}} / \det \mathbf{Q}$. This proof calculates this value by carrying out several decompositions of $\mathbf{Q}$ and $\widetilde{\mathbf{Q}}$. First, we will show step-by-step how $\det \mathbf{Q}$ is calculated. The same steps are then used to calculate $\det \widetilde{\mathbf{Q}}$ and therefore v_0.

Step 1: *Factor out terms involving links within module* 1.

First, write $\mathbf{Q}$ in block-matrix form $\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}$, where $\mathbf{A} = \begin{bmatrix} 1 & -p_1(n_1-1) \\ -p_1 & 1-p_1(n_1-2) \end{bmatrix}$ captures the terms describing links within module 1. The submatrices $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$ are defined accordingly. Then by the block-matrix determinant formula, $\det \mathbf{Q} = \det(\mathbf{A}) \det\left(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}\right)$.

Step 2: *Factor out terms involving the remaining modules in cluster 1, that is, modules* 2 *to* k_1.

We can write $\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$ in block matrix form as $\begin{bmatrix} \mathbf{M} & \mathbf{N} \\ \mathbf{O} & \mathbf{P} \end{bmatrix}$, where $\mathbf{M}$ captures the terms describing links within cluster 1, and each of these submatrices is defined as follows.

$$\boldsymbol{D} - \boldsymbol{C}\boldsymbol{A}^{-1}\boldsymbol{B} = \left[\begin{array}{ccc|ccccccc}
x_2^{-1} - \lambda^A n_2 & \cdots & -\lambda^A n_{k_1} & -\lambda^B n_{k_1+1} & \cdots & -\lambda^B n_{k_2} & \cdots & -\lambda^B n_{k_{K-1}+1} & \cdots & -\lambda^B n_{k_K} \\
\vdots & \boldsymbol{M} & \vdots & \vdots & \ddots & \vdots & \boldsymbol{N} & \vdots & \ddots & \vdots \\
-\lambda^A n_2 & \cdots & x_{k_1}^{-1} - \lambda^A n_{k_1} & -\lambda^B n_{k_1+1} & \cdots & -\lambda^B n_{k_2} & \cdots & -\lambda^B n_{k_{K-1}+1} & \cdots & -\lambda^B n_{k_K} \\
\hline
-\lambda^B n_2 & \cdots & -\lambda^B n_{k_1} & x_{k_1+1}^{-1} - \lambda^2 n_{k_1+1} & \cdots & -\lambda^2 n_{k_2} & \cdots & -\lambda n_{k_{K-1}+1} & \cdots & -\lambda n_{k_K} \\
\vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\
-\lambda^B n_2 & \cdots & -\lambda^B n_{k_1} & -\lambda^2 n_{k_1+1} & \cdots & x_{k_2}^{-1} - \lambda^2 n_{k_2} & \cdots & -\lambda n_{k_{K-1}+1} & \cdots & -\lambda n_{k_K} \\
\vdots & \boldsymbol{O} & \vdots & \vdots & \ddots & \vdots & \boldsymbol{P} & \vdots & \ddots & \vdots \\
-\lambda^B n_2 & \cdots & -\lambda^B n_{k_1} & -\lambda n_{k_1+1} & \cdots & -\lambda n_{k_2} & \cdots & x_{k_{K+1}+1}^{-1} - \lambda^K n_{k_{K-1}+1} & \cdots & -\lambda^K n_{k_K} \\
\vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \ddots & \vdots \\
-\lambda^B n_2 & \cdots & -\lambda^B n_{k_1} & -\lambda n_{k_1+1} & \cdots & -\lambda n_{k_2} & \cdots & -\lambda^K n_{k_{K-1}+1} & \cdots & x_{k_K}^{-1} - \lambda^K n_{k_K}
\end{array}\right]$$

where $\lambda = \frac{p}{1-pn_1x_1}$, $\lambda^\ell = \left(p^\ell - p\right) + \lambda$, $\lambda^A = \frac{p^1}{p}\left(1 + \left(p^1 - p\right) n_1 x_1\right)\lambda$, and $\lambda^B = \left(1 + \left(p^1 - p\right) n_1 x_1\right)\lambda$.

Again, by the block-matrix determinant formula, $\det\left(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}\right) = \det\left(\mathbf{M} - \mathbf{N}\mathbf{P}^{-1}\mathbf{O}\right)\det(\mathbf{P})$.

Step 3: *Write the terms involving clusters* 2 to K *as the sum of a diagonal matrix and a low-rank matrix.*

Next, note that we can write $\mathbf{P} = \mathbf{X}^{-1} + \mathbf{U}\mathbf{V}^T$, where

$$\mathbf{X} = \begin{bmatrix} x_{k_1+1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & x_{k_K} \end{bmatrix}, \mathbf{U} = \begin{bmatrix} -\lambda^2 & \cdots & -\lambda \\ \vdots & \ddots & \vdots \\ -\lambda^2 & \cdots & -\lambda \\ \vdots & \ddots & \vdots \\ -\lambda & \cdots & -\lambda^K \\ \vdots & \ddots & \vdots \\ -\lambda & \cdots & -\lambda^K \end{bmatrix}, \mathbf{V} = \begin{bmatrix} n_{k_1+1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ n_{k_2} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & n_{k_{K-1}} \\ \vdots & \ddots & \vdots \\ 0 & \cdots & n_{k_K} \end{bmatrix}.$$

Using the Weinstein-Aronszajn identity, $\det \mathbf{P} = \det\left(\mathbf{X}^{-1}\right)\det\left(\mathbf{I}+\mathbf{V}^T\mathbf{X}\mathbf{U}\right)$. The second term in this identity can, in turn, be written as the sum of a diagonal matrix and a rank-one matrix: $\mathbf{I}+\mathbf{V}^T\mathbf{X}\mathbf{U} = \mathbf{E}+\mathbf{u}\mathbf{v}^T$, where

$$\mathbf{E} = \begin{bmatrix} 1-\left(\lambda^2-\lambda\right)S^2 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1-\left(\lambda^K-\lambda\right)S^K \end{bmatrix}, \mathbf{u} = \begin{bmatrix} S^2 \\ \vdots \\ S^K \end{bmatrix}, \mathbf{v} = \begin{bmatrix} -\lambda \\ \vdots \\ -\lambda \end{bmatrix},$$

where recall that $S^k = \sum_{m\in\mathcal{M}^*,k=k(m)} n_m x_m$. By the matrix determinant lemma, $\det\left(\mathbf{I}+\mathbf{V}^T\mathbf{X}\mathbf{U}\right) = \left(1+\mathbf{v}^T\mathbf{E}^{-1}\mathbf{u}\right)\det\mathbf{E}$.

Step 4: *Rewrite the terms linking cluster* 1 *to clusters* 2 *to* K *in terms of calculable matrices.*

Recall from step 2 that calculating $\det\left(\mathbf{D}-\mathbf{C}\mathbf{A}^{-1}\mathbf{B}\right)$ requires calculating $\det\left(\mathbf{M}-\mathbf{N}\mathbf{P}^{-1}\mathbf{O}\right)$ and $\det\left(\mathbf{P}\right)$. Step 3 carried out the latter. This step shows how to calculate $\det\left(\mathbf{M}-\mathbf{N}\mathbf{P}^{-1}\mathbf{O}\right)$. Since $\mathbf{P} = \mathbf{X}^{-1}+\mathbf{U}\mathbf{V}^T$ is the sum of a diagonal matrix and a low-rank matrix, the Woodbury matrix identity allows us to write its inverse as follows

$$\mathbf{P}^{-1} = \mathbf{X} - \mathbf{X}\mathbf{U}\left(\mathbf{I}+\mathbf{V}^T\mathbf{X}\mathbf{U}\right)^{-1}\mathbf{V}^T\mathbf{X}.$$

Moreover, since $\mathbf{I}+\mathbf{V}^T\mathbf{X}\mathbf{U} = \mathbf{E}+\mathbf{u}\mathbf{v}^T$ is the sum of a diagonal matrix and a rank-one matrix, the Sherman-Morrison identity allows us to write its inverse as

$$\left(\mathbf{I}+\mathbf{V}^T\mathbf{X}\mathbf{U}\right)^{-1} = \mathbf{E}^{-1} - \frac{\mathbf{E}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{E}^{-1}}{1+\mathbf{v}^T\mathbf{E}^{-1}\mathbf{u}}.$$

Step 5: *Substitute in the expressions from steps 1 to 4 to give an expression for* $\det\mathbf{Q}$.

Putting together each of the preceding steps,we have the following:

$$\begin{aligned}
\det\mathbf{Q} &= \det\left(\mathbf{A}\right)\det\left(\mathbf{D}-\mathbf{C}\mathbf{A}^{-1}\mathbf{B}\right) \\
&= \det\left(\mathbf{A}\right)\det\left(\mathbf{M}-\mathbf{N}\mathbf{P}^{-1}\mathbf{O}\right)\det\mathbf{P} \\
&= \det\left(\mathbf{A}\right)\det\left(\mathbf{M}-\mathbf{N}\mathbf{P}^{-1}\mathbf{O}\right)\det\left(\mathbf{X}^{-1}\right)\det\left(\mathbf{I}+\mathbf{V}^T\mathbf{X}\mathbf{U}\right) \\
&= \det\left(\mathbf{A}\right)\det\left(\mathbf{M}-\mathbf{N}\mathbf{P}^{-1}\mathbf{O}\right)\det\left(\mathbf{X}^{-1}\right)\left(1+\mathbf{v}^T\mathbf{E}^{-1}\mathbf{u}\right)\det\mathbf{E} \\
&= \det\left(\mathbf{A}\right)\det\left(\mathbf{M}-\mathbf{N}\left(\mathbf{X}-\mathbf{X}\mathbf{U}\left(\mathbf{E}^{-1}-\frac{\mathbf{E}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{E}^{-1}}{1+\mathbf{v}^T\mathbf{E}^{-1}\mathbf{u}}\right)\mathbf{V}^T\mathbf{X}\right)\mathbf{O}\right)\cdot \\
&\quad \det\left(\mathbf{X}^{-1}\right)\left(1+\mathbf{v}^T\mathbf{E}^{-1}\mathbf{u}\right)\det\mathbf{E}.
\end{aligned}$$

Step 6: *Carry out the preceding steps for the matrix* $\widetilde{\mathbf{Q}}$ *to obtain an expression for* $\det\widetilde{\mathbf{Q}}$.

Recall that $\widetilde{\mathbf{Q}}$ is the matrix obtained by removing the first row and column of $\mathbf{Q}$. Carrying out steps 1 to 5 on this matrix, denote the corresponding matrices with tildes. For the analog of step 1, let $\widetilde{\mathbf{A}} = 1 - p_1 (n_1 - 2)$ and define $\widetilde{\mathbf{B}}$, $\widetilde{\mathbf{C}}$, and $\widetilde{\mathbf{D}}$ correspondingly. For the analog of step 2, write $\widetilde{\mathbf{D}} - \widetilde{\mathbf{C}}\widetilde{\mathbf{A}}^{-1}\widetilde{\mathbf{B}}$ in block matrix form as $\begin{bmatrix} \widetilde{\mathbf{M}} & \widetilde{\mathbf{N}} \\ \widetilde{\mathbf{O}} & \widetilde{\mathbf{P}} \end{bmatrix}$, where each of the submatrices $\widetilde{\mathbf{M}}$, $\widetilde{\mathbf{N}}$, $\widetilde{\mathbf{O}}$, and $\widetilde{\mathbf{P}}$ are the same as $\mathbf{M}$, $\mathbf{N}$, $\mathbf{O}$, and $\mathbf{P}$, except that each of the λ, λ^ℓ, λ^A, and λ^B terms are replaced with $\tilde{\lambda} = \frac{p}{1-p(n_1-1)\tilde{x}_1}$, $\tilde{\lambda}^\ell = \left(p^\ell - p\right) + \tilde{\lambda}$, $\tilde{\lambda}^A = \frac{p^1}{p}\left(1 + \left(p^1 - p\right)(n_1 - 1)\tilde{x}_1\right)\tilde{\lambda}$, and $\tilde{\lambda}^B = \left(1 + \left(p^1 - p\right)(n_1 - 1)\tilde{x}_1\right)\tilde{\lambda}$, where $\tilde{x}_1 = \frac{1}{1-p_1(n_1-2)+p(n_1-1)}$. For the analog of step 3, write $\widetilde{\mathbf{P}} = \mathbf{X}^{-1} + \widetilde{\mathbf{U}}\mathbf{V}^T$, where $\widetilde{\mathbf{U}}$ is the same as $\mathbf{U}$ but with λ and λ^ℓ replaced by $\tilde{\lambda}$ and $\tilde{\lambda}^\ell$. Additionally, $\mathbf{I} + \mathbf{V}^T\mathbf{X}\widetilde{\mathbf{U}} = \mathbf{E} + \mathbf{u}\tilde{\mathbf{v}}^T$, where $\tilde{\mathbf{v}}$ is the same as $\mathbf{v}$ but with λ replaced by $\tilde{\lambda}$.

Putting together each of these steps, we have

$$\begin{aligned}\det\widetilde{\mathbf{Q}} &= \det\left(\widetilde{\mathbf{A}}\right)\det\left(\widetilde{\mathbf{M}} - \widetilde{\mathbf{N}}\left(\mathbf{X} - \mathbf{X}\widetilde{\mathbf{U}}\left(\mathbf{E}^{-1} - \frac{\mathbf{E}^{-1}\mathbf{u}\tilde{\mathbf{v}}^T\mathbf{E}^{-1}}{1 + \tilde{\mathbf{v}}^T\mathbf{E}^{-1}\mathbf{u}}\right)\mathbf{V}^T\mathbf{X}\right)\widetilde{\mathbf{O}}\right)\cdot \\ &\quad \det\left(\mathbf{X}^{-1}\right)\left(1 + \tilde{\mathbf{v}}^T\mathbf{E}^{-1}\mathbf{u}\right)\det\mathbf{E}\end{aligned}$$

and therefore

$$\omega_{11}(\mathbf{C}_1) = v_0 = \frac{\det\widetilde{\mathbf{Q}}}{\det\mathbf{Q}} = \frac{\det\left(\widetilde{\mathbf{A}}\right)}{\det(\mathbf{A})}\frac{\det\left(\widetilde{\mathbf{M}} - \widetilde{\mathbf{N}}\left(\mathbf{X} - \mathbf{X}\widetilde{\mathbf{U}}\left(\mathbf{E}^{-1} - \frac{\mathbf{E}^{-1}\mathbf{u}\tilde{\mathbf{v}}^T\mathbf{E}^{-1}}{1+\tilde{\mathbf{v}}^T\mathbf{E}^{-1}\mathbf{u}}\right)\mathbf{V}^T\mathbf{X}\right)\widetilde{\mathbf{O}}\right)}{\det\left(\mathbf{M} - \mathbf{N}\left(\mathbf{X} - \mathbf{X}\mathbf{U}\left(\mathbf{E}^{-1} - \frac{\mathbf{E}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{E}^{-1}}{1+\mathbf{v}^T\mathbf{E}^{-1}\mathbf{u}}\right)\mathbf{V}^T\mathbf{X}\right)\mathbf{O}\right)}\frac{1 + \tilde{\mathbf{v}}^T\mathbf{E}^{-1}\mathbf{u}}{1 + \mathbf{v}^T\mathbf{E}^{-1}\mathbf{u}}.$$

Making the appropriate substitutions, we have

$$\omega_{11}(\mathbf{C}_1) = \frac{1 + \left(p_1 - p^1\right)x_1}{1 + p_1} + \frac{\left(p^1 - p\right)x_1^2}{1 - \left(p^1 - p\right)S^1} + \left(\frac{1}{1 - \left(p^1 - p\right)S^1}\right)^2\frac{px_1^2}{1 - p\sum_{k=1}^K \frac{S^k}{1-\left(p^k - p\right)S^k}},$$

which completes the proof. ■

PROPOSITION 6. *There exist thresholds $\lambda_i^k \geq 0$ such that it is optimal for agent $i \in \mathcal{N}$ to tell his state to a different agent $j \in \mathcal{N}$ who belongs to cluster $k \in \mathcal{K}$ if and only if:*

(i.) agent j belongs to the same module $m(j) = m(i)$ or

(ii.) agent j belongs to a different module $m(j) \neq m(i)$ with cohesion $x_{m(j)} \geq \lambda_i^k$.

Proof of Proposition 6. Let $\tilde{S}^1 = S^1 - n_1x_1$ and define the function $h\left(\tilde{S}^1, S^2, S^3, \ldots, S^K\right)$ to be

$$\frac{\left(p^1 - p\right)x_1^2}{1 - \left(p^1 - p\right)\left(n_1x_1 + \tilde{S}^1\right)} + \frac{1}{\left(1 - \left(p^1 - p\right)\left(n_1x_1 + \tilde{S}^1\right)\right)^2}\frac{px_1^2}{1 - p\left(\frac{\left(n_1x_1 + \tilde{S}^1\right)}{1-\left(p^1-p\right)\left(n_1x_1+\tilde{S}^1\right)} + \sum_{k=2}^K \frac{S^k}{1-\left(p^k-p\right)S^k}\right)}.$$

We will establish that this function is convex in each of its arguments and use this fact to argue that optimal communication networks have the multi-threshold property described in the lemma.

First, we show that h is convex in $\tilde{S}^1$. Let $W = \sum_{k=2}^{K} \frac{S^k}{1-\left(p^k-p\right)S^k}$. Then h is convex in $\tilde{S}^1$ if and only if

$$\tilde{h}\left(\tilde{S}^1\right) = \frac{\left(p^1-p\right)x_1^2}{1-\left(p^1-p\right)\left(n_1x_1+\tilde{S}^1\right)} + \frac{1}{\left(1-\left(p^1-p\right)\left(n_1x_1+\tilde{S}^1\right)\right)^2}\frac{px_1^2}{1-p\left(\frac{\left(n_1x_1+\tilde{S}^1\right)}{1-\left(p^1-p\right)\left(n_1x_1+\tilde{S}^1\right)}+W\right)}$$

is convex in $\tilde{S}^1$. This function is convex because it is twice differentiable and

$$\tilde{h}''\left(\tilde{S}^1\right) = 2x_1^2\left(\frac{p^1}{1-\left(p^1-p\right)S^1} + \frac{p}{1-\left(p^1-p\right)S^1}\frac{p^1\frac{S^1}{1-\left(p^1-p\right)S^1}+pW}{1-p^1\frac{S^1}{1-\left(p^1-p\right)S^1}-pW}\right)^3 > 0,$$

since $1-\left(p^1-p\right)S^1 > 0$ and $1-p^1\frac{S^1}{1-\left(p^1-p\right)S^1}-pW > 0$.

Next, we show that h is convex in S^ℓ for $\ell > 1$. Let $W = \frac{\left(n_1x_1+\tilde{S}^1\right)}{1-\left(p^1-p\right)\left(n_1x_1+\tilde{S}^1\right)} + \sum_{k\neq 1,k\neq \ell} \frac{S^k}{1-\left(p^k-p\right)S^k}$. Then h is convex in S^ℓ if and only

$$\tilde{h}\left(S^\ell\right) = \frac{1}{1-p\left(W+\frac{S^\ell}{1-\left(p^\ell-p\right)S^\ell}\right)}$$

is convex in S^ℓ. This function is convex because it is twice differentiable and

$$\tilde{h}''\left(S^\ell\right) = 2\frac{\tilde{h}\left(S^\ell\right)^3}{\left(1-\left(p^\ell-p\right)S^\ell\right)^3}p\left(p^\ell\left(1-pW\right)+\left(p\right)^2W\right) > 0,$$

since $1-pW > 0$ and $1-\left(p^\ell-p\right)S^\ell > 0$.

To establish the multi-threshold property, let $\mathcal{K}$ be an arbitrary set of modules with $m(i) \in \mathcal{K}$. Define $S^\ell(\mathcal{K}) = \sum_{m:m\in\mathcal{K},k(m)=\ell} n_m x_m$, and denote by $\mathbf{C}_i(\mathcal{K})$ the row of the communication matrix in which agent i tells θ_i to agent j if and only if $m(j) \in \mathcal{K}$. By Lemma D1, $R_i(\mathbf{C}_i(\mathcal{K}))$ is a linear function of $h\left(\tilde{S}^1(\mathcal{K}), S^2(\mathcal{K}), \ldots, S^K(\mathcal{K})\right)$.

Now suppose that it is optimal to inform all modules in $\mathcal{K}$. Then it must be the case that for all $m \in \mathcal{K}$ and $k(m) = \ell$,

$$\begin{aligned}\gamma &\leq \frac{h\left(\tilde{S}^1(\mathcal{K}), S^2(\mathcal{K}), \ldots, S^K(\mathcal{K})\right) - h\left(\tilde{S}^1(\mathcal{K}), S^2(\mathcal{K}), \ldots, S^\ell(\mathcal{K}\backslash\{m\}), \ldots, S^K(\mathcal{K})\right)}{n_m} \\ &< x_m h_\ell\left(\tilde{S}^1(\mathcal{K}), S^2(\mathcal{K}), \ldots, S^K(\mathcal{K})\right),\end{aligned}$$

where the second inequality holds because h is convex in $S^\ell(\mathcal{K})$ and where h_ℓ denotes the partial derivative of h with respect to S^ℓ. Suppose further that it is not optimal to also inform some module $m' \notin \mathcal{K}$ with $k(m') = \ell$. Then it must be the case that

$$\begin{aligned}\gamma &> \frac{h\left(\tilde{S}^1(\mathcal{K}), S^2(\mathcal{K}), \ldots, S^\ell(\mathcal{K}\cup\{m'\}), \ldots, S^K(\mathcal{K})\right) - h\left(\tilde{S}^1(\mathcal{K}), S^2(\mathcal{K}), \ldots, S^K(\mathcal{K})\right)}{n_{m'}} \\ &> x_{m'} h_\ell\left(\tilde{S}^1(\mathcal{K}), S^2(\mathcal{K}), \ldots, S^K(\mathcal{K})\right).\end{aligned}$$

These two inequalities imply that $x_m > x_{m'}$ for all modules m in cluster ℓ that are are optimally told about θ_i and modules m' in cluster ℓ that are optimally not told about θ_i. In other words, there is a threshold λ_i^ℓ such that agent i tells θ_i to agent j in module $m(j) \neq m(i)$ and cluster $k(m(j)) = \ell$ if and only if $x_{m(j)} \geq \lambda_i^\ell$. ■

We conclude this appendix by establishing that the principal's problem is supermodular for arbitrary production networks satisfying $p_{ii} = 0$, $p_{ij} = p_{ji}$, *and* $\sum_{j=1}^N p_{ij} < 1$.

PROPOSITION 7. *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 7. 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 $\omega_{ii}(\boldsymbol{C}_i)$ 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 $\left(a_i^2\sigma_i^2, \{p_{ij}\}_{ij}, \{c_{ij}\}_{ij}, -\gamma\right)$, so the comparative statics results follow from Topkis's theorem. It remains, therefore, to show that $\omega_{ii}(\boldsymbol{C}_i)$ is supermodular in $\boldsymbol{C}_i$.

To show that $\omega_{ii}(\cdot)$ is supermodular, let $\mathcal{J} \subset \mathcal{N}$ denote a subset of agents, and denote by $\mathbf{c}(\mathcal{J})$ 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. ■

Appendix E: Incentive Conflicts

PROPOSITION 8. *If agents internalize only a fraction $\rho \in [0,1]$ of the needs to coordinate, an optimal communication network solves*

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

where

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

and

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

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

Proof of Proposition 8. Suppose agents internalize only fraction $\rho \in [0,1]$ of the need to coordinate. Then they act as if the need for coordination between d_i and d_j is ρp_{ij} rather than p_{ij}. Agent i's best response function is therefore

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

and so by Lemma 1, equilibrium decisions are given by

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

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

Given equilibrium decision-making, expected revenue can be written as

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

since by the law of iterated expectations, $\mathrm{E}\left[d_i^*\left(\mathrm{E}_i\left[d_j^*\right]-d_j^*\right)\right]=\mathrm{E}\left[\mathrm{E}_i\left[d_i^*\right]\left(\mathrm{E}_i\left[d_j^*\right]-\mathrm{E}_i\left[d_j^*\right]\right)\right]=0$, and $\mathrm{E}\left[d_i^*\mathrm{E}_i\left[d_j^*\right]\right]=\mathrm{E}\left[d_i^*d_j^*\right]$. Moreover, $\mathrm{E}\left[d_i^*\right]=0$, so $\mathrm{E}\left[d_i^*\theta_i\right]=\mathrm{Cov}\left(d_i^*,\theta_i\right)$ and $\mathrm{E}\left[d_i^*d_j^*\right]=\mathrm{Cov}\left(d_i^*,d_j^*\right)$. Next, note that since $\mathrm{Cov}\left(\theta_j,\theta_i\right)=0$ for all $i\neq j$,

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

and

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

which establishes the result. ■