Organizing Modular Production*

Niko Matouschek
Northwestern University

Michael Powell
Northwestern University

Bryony Reich
Northwestern University

Preliminary and Incomplete–Please Do Not Distribute
August 31, 2021

Abstract

We characterize the optimal communication network in a firm with a modular production function, which we model as a network of decisions with a non-overlapping community structure. Optimal communication is characterized by two nested hierarchies that determine who each agent receives information from and sends information to. Receiver rank depends only on the cohesion of an agent's module while sender rank also depends on the importance of adaptation. We provide conditions for the Mirroring Hypothesis and show that optimal communication is consistent with a core-periphery structure in which the core is formed by the most cohesive modules.

*We thank Hossein Alidaee for excellent research assistance. Matouschek: n-matouschek@kellogg.northwestern.edu; Powell: mike-powell@kellogg.northwestern.edu; Reich: bryony.reich@kellogg.northwestern.edu

1 Introduction

Almost 60 years ago, Herbert Simon observed that complex social, technological, and biological systems–large firms, mechanical watches, the brain–tend to be made up of communities or modules, groups of elements with stronger within than across group interactions (Simon 1962). The advantage of this modular structure, he argued, is that it allows systems to adapt to changes in the environment by adjusting a small number of modules while leaving the rest of the system unchanged. The prevalence of modular structures has since been confirmed by the flourishing literature on community detection, which has documented them in a wide variety of contexts from the internet to the global air transportation network and the human brain (Meunier et al. 2009, Guimera et al. 2005, and Fortunato 2009) .

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

The move towards modular production has not been confined to the computer industry. Over the last few decades, firms across a wide range of industries followed in IBMs footsteps and developed products with modular production functions. Smart phones, tablets, electric cars–the defining products of our times–are all made by assembling a limited number of modules. Even homes are now routinely assembled from pre-made modules rather than built on site from scratch. Nor is this move towards modular production confined to physical products. Modular programming–separating the different functions of computer programs into independent and interchangeable modules–was developed in the 1960s and is now a common feature of almost all programming languages. Of course, many products exhibited some degree of modularity even before the System/360. Builders installed pre-made doors and windows long before the rise of modular home building. What is different now, though, is that many products are modular by design. They are produced entirely by assembling a limited number of modules and, in line with Herbert Simon's observations, they are now more rule than exception.

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

The goal of this paper is to take a first step towards exploring the economic implications of modular production ("first step" in the economics literature; as discussed below there is an expansive existing literature in management). We focus on the internal organization of a single firm with a modular production function, which we model as a network of decisions with a non-overlapping community structure. Each decision is made by a different agent who tries to adapt it to his local information while also coordinating it with the decisions made by the other agents. The decisions are partitioned into "modules," groups of decisions that require more coordination with each other than with decisions in other modules. The adjacency matrix of the production network, therefore, takes the form of a block matrix, as illustrated in Figure 1a. This structure approximates the "design structure matrix" of a modular laptop computer, which we reproduce from McCord and Eppinger (1993) in Figure 1b.

The only impediment to efficient production is that agents involved in the making of one module do not observe the information relevant for the production of other modules. To improve efficiency, the firm can establish communication channels between agents who work on different modules. An IBM engineer who works on the disk drive does not directly observe the factors relevant to someone who works on the processor. But IBM can require the former to meet with the latter and learn about his local information. The problem is that such communication does not come for free. Even in the age of ever-evolving communication technologies, explaining the issues one faces, and understanding those faced by others, takes time and energy. Given this trade-off between the efficiency of decision-making and the cost of communication, what is the optimal communication network? In a firm with a modular production function, who should report to whom?[1]

[1] Our focus on the trade-off between the efficiency of decision-making and the cost of communication is in line with Kenneth Arrow's discussion of the design of optimal communication within firms in Arrow (1974), where he observes that: "*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*

The challenge in designing an optimal communication network is the abundance of possibilities and absence of an apparent way to order them. To see how the firm can overcome this challenge, it is useful to start by asking when it should add a single, directed link to an existing communication network. The cost of establishing such a link is the time and energy it takes the sender to explain his information to the receiver, which we take to be exogenous. The benefit is that learning the sender's information allows the receiver to coordinate his decision more closely with the sender's, which, in turn, allows the sender to adapt his decision more closely to his state. Crucially, this benefit is independent of what the receiver, or any other agent, knows about any other state. Because of this independence, the problem of designing an optimal communication network can be broken into a number of independent subproblems. It is sufficient for the firm to consider each agent in turn and ask who this agent should inform about his state. This separability result is what allows us to characterize optimal communication networks and establish our central result.

Our central result is that modular production gives rise to hierarchies. We fully characterize the optimal communication network and show that it consists of two nested hierarchies, one that determines who each agent sends his information to and another that determines who he receives information from. An agent with a higher sender rank sends his information to any agent that a lower ranked agent sends his information to. And an agent with a higher receiver rank receives information from any agent that a lower ranked agent receives information from.

An agent's rank in either hierarchy depends critically on the cohesiveness of his module, which is increasing in the number of decisions that are part of the module and the need for coordination between them and decreasing in the "degree of coupling," the need for coordination across modules. The more cohesive an agent's module is, the more important it is that the agent learns about decisions in other modules and that agents working on other modules learn about his.

Receiver rank is fully determined by module cohesion. Agents working on the same module have the same rank, they either all learn about a decision in another module or none of them do. The same is not true for sender rank, which can vary across agents working on the same module. The reason is that the benefit of sharing information does not only depend on module cohesion but also on the variability of the sender's information and the need to adapt his decision to it. If an agent's module is very cohesive, it is important he learns about other modules to enable them to adapt their decisions to their information. But if his own information is predictable, or the need to adapt his decision to his information is low, it may not be important for those working on other modules to learn his information. Sender and receiver ranks need not be perfectly correlated.

channels. The optimal choice of internal communication structures is a vastly difficult question."

Agents who hear the most may differ from those who speak the most.

We apply our characterization to explore the notion that communication links should simply mirror technological interdependency, that "*we should expect to see a very close relationship...between a network graph of technical dependencies within a complex system and network graphs of organizational ties showing communication channels*" (Colfer and Baldwin 2016, p.713) This notion has a long history in a number of fields, including organizational and product-design, where it is known as the "Mirroring Hypothesis," and computer science, where it is referred to as "Conway's Law" (Thompson (1967), Conway (1967) and, for a discussion of the literature, Colfer and Baldwin (2016)). In our context, the Mirroring Hypothesis predicts communication within but not across modules. We show that the Mirroring Hypothesis holds, provided that the organization is moderate. It cannot consist of too many modules, and none of them can have too many members are require too much coordination among them. If it did, the organization benefits from communication across modules. At the same time, none of the modules can have too few members and require too little coordination among them, which would negate the need for communication even within modules.

Our characterization result also speaks to the possibility that firms engage in "partial mirroring" by limiting across-module communication to clusters of modules. The result implies that if such clustering is optimal, there is at most one cluster of modules whose agents engage in across-module communication. Moreover, this cluster is made up of the most cohesive modules. Agents working on less cohesive modules are outside of the cluster and do not communicate across modules. In this case, the optimal communication network, therefore, has a core-periphery structure, which is a pervasive structure among social and communication networks (Herskovic and Ramos 2020). The model shows that this common structure can be an optimal way to organize a modular firm, but only if the core is made up of the modules with the largest number of decisions and/or the highest need for coordination among them.

2 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 $d_i \in \mathbb{R}$ that is associated with a state $\theta_i \in \mathbb{R}$, where $\mathcal{N} = \{1, ..., N\}$ is the set of agents. Output depends on how well each decision is adapted to

its associated state and coordinated with the other decisions. Specifically, output is given by

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

where $a_i>0$ captures the importance of adapting decision d_i to its state θ_i, $p_{ij}\geq 0$ represents the need to coordinate decisions $d_i\neq d_j$, and where $p_{ii}=0$. The need for coordination is bidirectional, that is, $p_{ij}=p_{ji}$. Production, therefore, takes the form of 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 optimal decisions exist. Finally, we normalize the price of the product to one so that output (1) also represents revenue.

Modules. Each agent-decision-state triplet (i,d_i,θ_i) belongs to a "module" $\mathcal{M}_m$ for $m\in\{1,...,M\}$, which is a set of $n_m\geq 1$ such triplets. We label agents and their associated decisions and states such that $\left(1,d_1,\theta_1\right),...,\left(n_1,d_{n_1},\theta_{n_1}\right)$ belong to module $\mathcal{M}_1$, $\left(n_1+1,d_{n_1+1},\theta_{n_1+1}\right),...,\left(n_1+n_2,d_{n_1+n_2},\theta_{n_1+n_2}\right)$ belong to module $\mathcal{M}_2$, and so on. The function $m\left(i\right)$ gives the module $\mathcal{M}_{m(i)}$ that (i,d_i,θ_i) belongs to.

The need for coordination is stronger between two decisions within the same module than across two modules. Specifically, the need for coordination between any two decisions d_i and $d_j\neq d_i$ is given by $p_{ij}=t\geq 0$ if they belong to different modules and by $p_{ij}=p_m\geq t$ if they belong to the same module $\mathcal{M}_m$.

Information & Communication. Each state θ_i is independently drawn from a distribution with zero mean and variance σ_i^2, where $i\in\mathcal{N}$. The realization of state θ_i is observed by agent i and the other agents in his module $\mathcal{M}_{m(i)}$. All other information is public.

After agent i has learnt the realization of his state θ_i he may communicate it to agents in other modules. If agent i communicates his state to agent j, agent j learns it perfectly, otherwise agent j learns nothing about θ_i. Moreover, communication is directional so that agent i may tell agent j about θ_i without agent j telling agent i about θ_j. Communication, therefore, takes the form of a directed network, which we summarize in the $N\times N$ matrix $\boldsymbol{C}$. Entry c_{ij} is equal to one if agent i informs agent j about his state, or agent j observes θ_i directly, and it is equal to zero otherwise. Row $\boldsymbol{C}_i$ then summarizes the agents who learn θ_i and column $\boldsymbol{C}_j$ summarizes the states agent j learns about.

Organizational Design. The communication network is designed by the principal. Each communication link involves a resource cost $\gamma>0$ that captures the resources necessary to engage in communication.

At the beginning of the game, the principal chooses an optimal communication network to maximize expected profits. The principal's problem is therefore given by

$$\max_{\boldsymbol{C}} \mathrm{E}[r\left(d_1, ..., d_N\right) | \boldsymbol{C}] - \gamma \sum_{i=1}^{N} \left(\boldsymbol{C}_i \mathbf{1} - n_{m(i)}\right). \tag{2}$$

Timing. After the principal designed the communication network, agents learn their states and communicate them to the agents specified in the network. Next, the agents simultaneously and independently make their decisions, payoffs are realized, and the game ends. The solution concept we use is Perfect Bayesian Equilibrium.

3 Decision Making

We solve the game by backward induction. After agents have observed their states and communicated with each other, they make the decisions that solve

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

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

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

Each agent's best-response is the weighted sum of his state and the decisions he expects the other agents to make, where the weight on his own state is a_i and the weight on the decision he expects agent j to make is p_{ij}. Solving this system of equations we get our first lemma.

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

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

where the jth row of the communication network $\boldsymbol{C}_j$ summarizes who knows state θ_j and where $\varpi_{ij}\left(\boldsymbol{P}, \boldsymbol{C}_j\right)$ denotes the ijth entry of $\left(\boldsymbol{I} - \left(\mathrm{diag}\boldsymbol{C}_j\right)\boldsymbol{P}\right)^{-1}$.

Agent i's optimal decision d_i^* is the weighted sum of all states, where the weight on state θ_j is given by a_j, which captures the need to adapt decision d_j to the state, times $\varpi_{ij}\left(\boldsymbol{P}, \boldsymbol{C}_j\right)$, which is the value of all walks from nodes i to j on the subgraph of the production network that consists

only of decisions made by agents who know state θ_j. If agent i does not know θ_s, for instance, d_i is not part of this subgraph, and so $\varpi_{is}(\boldsymbol{P}, \boldsymbol{C}_j) = 0$. 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 reflects a general implication of the lemma that will be important for what follows: for a given production network, the weight agent i puts on state θ_s depends only on who knows θ_s and not on what agent i, or any of the other agents, know about any other state.

To get an intuition for decision rule (4), take the production network $\boldsymbol{P}^e$ in Figure 2 and consider how much weight agent 1 puts on his own state. If only agent 1 observes his state, agents 2 and 3 cannot put any weight on θ_1. As a result, agent 1 faces a trade-off between adapting his decision to θ_1 and coordinating it with d_2 and d_3. Decision rule (4) shows that if agent 1 resolves this trade-off optimally, the weight he puts on his own state is given by

$$a_1 \varpi_{11}(\boldsymbol{P}^e, (1,0,0)) = a_1,$$

where the equality follows from the fact that the value of all walks from d_1 to d_1 that go only through d_1 is 1.

If agent 2 learns θ_1, his decision will put some weight on this state, which relaxes agent 1's trade-off between adapting his decision to θ_1 and coordinating it with d_2. As a result, agent 1 increases the weight he puts on his state to

$$a_1 \varpi_{11}(\boldsymbol{P}^e, (1,1,0)) = a_1 \left(1 + \frac{t^2}{1-t^2}\right),$$

where the fraction on the RHS is the value of all walks from d_1 back to d_1 that go through d_1 and d_2.

Finally, if agent 3 also learns θ_1, the weight agent 1 puts on his state increases by even more to

$$a_1 \varpi_{11}(\boldsymbol{P}^e, (1,1,1)) = a_1 \left(1 + \frac{t^2}{1-t^2} + \frac{t^2}{1-t^2} + \frac{2t^2\left(t^2 + p_{23}\right)}{\left(1-t^2\right)\left(1-p_{23}-2t^2\right)}\right),$$

where the second fraction is the value of all walks from d_1 back to d_1 that only go through d_1 and d_3 and the third fraction is the value of all walks from d_1 back to d_1 that go through all three decisions. The intuition for the second fraction is the same as for the first: if agent 3 learns θ_1, his decision will put some weight on this state. This relaxes the trade-off agent 1 faces between adapting d_1 to θ_1 and coordinating it with d_3, inducing him to put more weight on his state. The additional effect that is captured by the third fraction above is that by making his decision vary

with θ_1, agent 3 also relaxes the trade-off agent 2 faces between coordinating d_2 with both d_1 and d_3. As a result, agent 2 increases the weight his decision puts on θ_1, which allows agent 1 to do the same.

The example illustrates three general features of the agents' decision-making that matter for what follows. First, if only agent i observes θ_i, he puts weight a_i on his state. The parameter a_i therefore captures the degree of autonomous adaptation. Second, allowing additional agents to observe θ_i multiplies the weight agent i puts on his state to $\varpi_{ii}(\boldsymbol{P}, \boldsymbol{C}_j) a_i \geq a_i$. It does so since allowing additional agents to observe θ_i enables them to coordinate their decisions with each other and with d_i which, in turn, allows agent i to adapt his decision to θ_i without sacrificing as much coordination. The total weight that agent i puts on his state is therefore the product of the "coordination multiplier" $\varpi_{ii}(\boldsymbol{P}, \boldsymbol{C}_j)$ and the degree of autonomous adaptation a_i. Finally, the increase in the weight agent i puts on his state if agent $j \neq i$ learns θ_i is larger, if agent $k \neq i, j$ knows θ_i. We summarize these properties in the following corollary.

COROLLARY. *The weight $a_i \varpi_{ii}(\boldsymbol{P}, \boldsymbol{C}_i)$ that agent i's decision d_i^* puts on his state θ_i satisfies $\varpi_{ii}(\boldsymbol{P}, \boldsymbol{I}_i) a_i = a_i$ and is increasing and supermodular in $\boldsymbol{C}_i$.*

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

4 The Principal's Problem

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

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

which, for optimal decisions, becomes

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

In the proof of the next lemma we show that the expected value of the the second term on the RHS is zero. As a result, expected revenue is fully determined by the covariance between each decision and its state.

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

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

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

Expected revenue, therefore, boils down to how well each decision is adapted to its associated state. The extent to which agent i adapts his decision to his state depends on the communication network and, in particular, on who knows his state. If an additional agent observes θ_i, agent i puts more weight on his state, as we saw above. And if agent i puts more weight on his state, his decision is more closely adapted to his state, as captured by the larger covariance between his decision and state.

Next we substitute expected revenue (6) into the principal's problem (2) to obtain our first proposition.

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

$$\max_{\boldsymbol{C}_i} \sigma_i^2 a_i^2 \varpi_{ii}\left(\boldsymbol{P}, \boldsymbol{C}_i\right) - \gamma\left(\boldsymbol{C}_i \mathbf{1} - n_{m(i)}\right) \quad \textit{for all } i \in \mathcal{N}. \tag{7}$$

The proposition shows that to determine an optimal communication network, the principal can consider each agent in isolation and ask who this agent should inform about his state. Since $\varpi_{ii}\left(\boldsymbol{P}, \boldsymbol{C}_i\right)$ depends on $\boldsymbol{C}_i$ but not $\boldsymbol{C}_{-i}$–on who knows θ_i but not on who knows the other states–the answer to who agent $i \in \mathcal{N}$ should inform about his state in independent of who any other agent should inform about his state. For the remainder of the paper we denote the first term in (7) by

$$R_i\left(\boldsymbol{C}_i\right) \equiv \sigma_i^2 a_i^2 \varpi_{ii}\left(\boldsymbol{P}, \boldsymbol{C}_i\right)$$

and refer to it as the expected revenue generated by agent i.

The separability result in Proposition 1 greatly facilitates the principal's search for an optimal communication network. Supermodularity of $\varpi_{ii}\left(\boldsymbol{P}, \boldsymbol{C}_i\right)$ further simplifies the problem since it implies that whenever it is optimal to inform an agent about a state, it must also be optimal to inform all the other agents in the same module. The principal's problem, therefore, reduces to which modules each agent should inform about his state.

To state the principal's problem in these terms, let $\boldsymbol{D}$ denote the matrix that specifies which modules are informed about each state and let $\boldsymbol{F}$ denote the conversion matrix that specifies what module each agent belongs to. Specifically, $\boldsymbol{D}$ is an $N \times M$ matrix in which entry d_{im} is equal to

one if agent i tells agents in module $\mathcal{M}_m$ about his state θ_i and d_{im} is equal to zero if he does not. Moreover, $d_{im(i)} = 1$, which reflects that agents observe the states in their own module without having to be told about them. The conversion matrix $\boldsymbol{F}$, in turn, is an $M \times N$ matrix in which entry f_{mi} is equal to one if agent i belongs to module $\mathcal{M}_m$ and zero otherwise. We can then restate the principal's problem (7) as

$$\max_{\boldsymbol{D}_i} R_i\left(\boldsymbol{D}_i \boldsymbol{F}\right) - \gamma \left(\sum_{m=1}^{M} d_{im} n_m - n_{m(i)} \right) \quad \text{for all } i \in \mathcal{N}. \tag{8}$$

Finally, supermodularity of $\varpi_{ii}(\cdot)$, together with the linearity of the communication costs, imply that the subproblems (8) 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 so that we can understand the properties of optimal communication networks for any parameter values. We tackle this problem in the next section.

5 Optimal Communication Networks

The separability result in Proposition 1 allows us to solve the principal's problem of designing an optimal communication network by considering each agent in isolation and asking who he should tell about his state. To economize on notation, we focus on agent 1, who is a member of module $\mathcal{M}_1$. Once we know who agent 1 should tell about his state we can apply answer to all the other agents and thus solve the principal's problem.

To understand who agent 1 should tell about his state, consider the expected revenue $R_1(\cdot)$ he generates if his state is known to the agents in his own module $\mathcal{M}_1$ and to those in some arbitrary set of other modules. Since the naming of modules is immaterial, there is no loss in denoting the modules whose agents know θ_1 by $\mathcal{M}_1$, ..., $\mathcal{M}_\ell$ for $\ell \in \{1, ..., M\}$. We can then define $\boldsymbol{D}_1(\ell)$ as the $1 \times M$ row vector that specifies the modules whose agents know θ_1 and $\boldsymbol{C}_1(\ell) = \boldsymbol{D}_1(\ell)\boldsymbol{F}$ as the corresponding $1 \times N$ row vector of the communication matrix that specifies the agents who know θ_1. The next lemma uses this notation to express the expected revenue agent 1 generates.

LEMMA 3. *Suppose agent 1's state θ_1 is known to all the agents in modules $\mathcal{M}_1$, ..., $\mathcal{M}_\ell$, for $\ell \in \{1, ..., M\}$, and none of the agents in other modules. Agent 1's expected revenue is then given by*

$$R_1\left(\boldsymbol{C}_1(\ell)\right) = a_1^2 \sigma_1^2 \left(\frac{1 - (n_1 - 2)\, p_1}{(1 + p_1)\left(1 - (n_1 - 1)\, p_1\right)} + \frac{t^2 x_1^2 \sum_{m=2}^{\ell} n_m x_m}{(1 - t n_1 x_1)\left(1 - t \sum_{m=1}^{\ell} n_m x_m\right)} \right), \tag{9}$$

where

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

The object x_m in the lemma captures the "cohesion" of module $\mathcal{M}_m$, which is increasing in its size and the need for coordination among its members and decreasing in the degree of coupling t. A module is, therefore, more cohesive if coordination among its members is relatively more important than coordination between its members and those in other modules.[2] Both module cohesion x_m and its scaled analog $n_m x_m$ play a key role in who agent 1 should inform.

The first fraction in the expression of agent 1's expected revenue (9) is the value of all walks from node 1 back to itself that only go through nodes in $\mathcal{M}_1$. Because these walks only go through module $\mathcal{M}_1$, their value depends only on its characteristics p_1 and n_1. The second fraction in (9), in turn, is the value of the additional walks that go through the other informed modules. Notice that the value of these additional walks depends on the characteristics of $\mathcal{M}_1$ only through x_1 and $n_1 x_1$ and it depends on the characteristics of the other informed modules only through the sum of their $n_m x_m$s. The revenue generated by agent 1, for instance, is the same whether he informs agents in one module with $n_2 x_2 = 10$ or agents in ten modules with $n_2 x_2 = ... = n_{11} x_{11} = 1$. This property of expected revenue is important for what follows.

Given the expression for expected revenue (9), we can determine when the principal benefits from having agent 1 inform additional agents about his state. Suppose that, initially, there are at least two modules whose agents do not know θ_1, that is, $\ell \leq M-3$, and that the principal expands the set of informed agents by having agent 1 also report his state to the agents in module $\mathcal{M}_{\ell+1}$. She benefits from doing so as long as the per node marginal revenue it generates is larger than the marginal cost, that is, as long as

$$\frac{1}{n_{\ell+1}} \left(R_1 \left(\boldsymbol{C}_1 \left(\ell + 1 \right) \right) - R_1 \left(\boldsymbol{C}_1 \left(\ell \right) \right) \right) \geq \gamma.$$

Using (9), the per node marginal revenue is given by

$$\frac{1}{n_{\ell+1}} \left(R_1 \left(\boldsymbol{C}_1 \left(\ell + 1 \right) \right) - R_1 \left(\boldsymbol{C}_1 \left(\ell \right) \right) \right) = \frac{t^2 a_1^2 \sigma_1^2 x_1^2 x_{\ell+1}}{\left(1 - t \sum_{m=1}^{\ell} n_m x_m \right) \left(1 - t \sum_{m=1}^{\ell+1} n_m x_m \right)}. \tag{10}$$

Notice that this expression is increasing in x_1 and $n_1 x_1$ and thus in n_1 and p_1. The more members module $\mathcal{M}_1$ has, or the more important it is to coordinate their decisions, the more the principal gains from agents in an additional module learning θ_1. Per node marginal revenue is

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

also increasing in a_1 and σ_1, the two parameters that capture potential differences across decisions within module $\mathcal{M}_1$. The larger either of them is, the more the principal gains from agent 1 adapting his decision to his state, and the more, thus, she benefits from having him communicate his state to other agents. Beyond the characteristics of agent 1's module and decision, marginal revenue (10) is increasing in the number of informed modules, which reflects the supermodularity of $\varpi_{11}(\cdot)$. And it is increasing in p_m and n_m for all $m = 2, ..., \ell + 1$, which reflects the fact that the value of the additional walks that are created by informing agents in another module about a state is increasing in the number of informed nodes and the need for coordination among them.

Suppose now that the principal extends the communication network by having agent 1 inform not just the agents in $\mathcal{M}_{\ell+1}$ but also those in $\mathcal{M}_{\ell+2}$. Again using (9), the per node marginal revenue of doing so

$$\frac{1}{n_{\ell+1}+n_{\ell+2}}\left(R_1\left(\boldsymbol{C}_1\left(\ell+2\right)\right)-R_1\left(\boldsymbol{C}_1\left(\ell\right)\right)\right)$$

is given by

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

Subtracting (10) from this expression, we have that the difference in the per node marginal revenues

$$\frac{1}{n_{\ell+1}+n_{\ell+2}}\left(R_1\left(\boldsymbol{C}_1\left(\ell+2\right)\right)-R_1\left(\boldsymbol{C}_1\left(\ell\right)\right)\right)-\frac{1}{n_{\ell+1}}\left(R_1\left(\boldsymbol{C}_1\left(\ell+1\right)\right)-R_1\left(\boldsymbol{C}_1\left(\ell\right)\right)\right)$$

is equal to

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

Notice that this expression is positive if $x_{\ell+2} \geq x_{\ell+1}$. If module $\mathcal{M}_{\ell+2}$ is more cohesive than $\mathcal{M}_{\ell+1}$, the per node marginal benefit from informing agents in both $\mathcal{M}_{\ell+1}$ and $\mathcal{M}_{\ell+2}$ is, therefore, larger than that of only informing only those in $\mathcal{M}_{\ell+1}$. It can then never be optimal for agent 1 to inform agents in $\mathcal{M}_{\ell+1}$ but not those in $\mathcal{M}_{\ell+2}$. It may be optimal for agent 1 to inform agents in both modules, neither, or only those in $\mathcal{M}_{\ell+2}$, but informing only the agents in $\mathcal{M}_{\ell+1}$ cannot be optimal. This, in turn, implies our main result: optimal communication follows a threshold rule in which agent 1, informs agents in any module $\mathcal{M}_m$ whose cohesion x_m is above a threshold and none of those whose value is below. Since there is nothing special about agent 1, the same applies to to any other agent. We can then state the solution to the principal's problem.

PROPOSITION 2. *Optimal communication is characterized by N thresholds $\lambda_i \geq 0$, one for each agent $i \in \mathcal{N}$. Given any two agents i and j who belong to different modules, agent i tells agent*

j about his state if and only if the cohesion of agent j's module is above agent i's threshold, that is, if and only if $x_{m(j)} \geq \lambda_i$. The threshold λ_i is increasing in marginal communication costs γ and decreasing in the need to adapt agent i's decision to his state $a_i\sigma_i$, the need to coordinate the decisions within his module $p_{m(i)}$, and the size of his module $n_{m(i)}$.

To illustrate the proposition for agent 1, it is convenient to label modules $\mathcal{M}_2, ..., \mathcal{M}_\ell$ in decreasing order of their values of x_m so that $\mathcal{M}_2$ denotes the module, other than $\mathcal{M}_1$, with the highest value of x_m and $\mathcal{M}_M$ denotes the one with the smallest. In Figure 3, the blue curve is the piecewise linear extension of expected revenue $R_1(\boldsymbol{C}_1(\ell))$, which we denote by $\overline{R}_1(\boldsymbol{C}_1(\ell))$, and the red line is a continuous representation of the costs $\sum_{m=1}^{\ell} n_m\gamma$. The changing curvature of expected revenue $\overline{R}_1(\boldsymbol{C}_1(\ell))$ reflects the countervailing economic forces at work. The supermodularity at the heart of the model pushes towards convexity while the community structure pushes towards concavity. A reduction in γ favors informing more modules because it flattens the cost curve. And an increase in a_1, σ_1^2, n_1, or p_1 favors informing more modules because it increases the per node marginal benefit (10) and thus steepens the expected revenue curve.

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

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

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

The fact that cohesion x_m depends only on the characteristics of module $\mathcal{M}_m$, and not on those of any other modules, also implies a sender hierarchy, in which a higher ranked agent informs all

the agents that a lower ranked agent informs, and possibly others. The rank of agent i, though, does not depend on $x_{m(i)}$ but on $\lambda_{m(i)}$, and it does so inversely. Agent i outranks agent j if $\lambda_{m(i)} \leq \lambda_{m(j)}$.

COROLLARY. *Optimal communication gives rise to a sender hierarchy among agents. For any agents $i, j, k \in \mathcal{N}$ who belong to different modules, if agent i's threshold λ_i is smaller than agent j's threshold λ_j, then agent j tells agent k about his state only if agent i also does.*

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

A distinct feature of both hierarchies is that they are nested. This is in contrast with the knowledge hierarchies in Garicano (2000), and the literature that builds on his work, in which an agent reports to his immediate boss but not to the boss's superiors. The choice of whether to inform the boss's superior is left with the boss, who sometimes decides to do so and sometimes does not. Such narrow communication chains are not optimal in our setting. If an agent informs a superior, he informs the superior's superiors. And if he hears from a subordinate, he hears from the subordinate's subordinates.

6 Application–The Mirroring Hypothesis

We now use the characterization of optimal communication networks in Proposition 2 to examine the mirroring hypothesis. As we discussed earlier, the mirroring hypothesis predicts that modular production favors modular organizations, organizations with intense communication within modules

but sparse communication across. To allow for within-module communication, we now relax the assumption that agents in the same module observe each others' states and assume instead that each state is observed only by the associated agent. An organization then mirrors the production function if the principal places communication links within modules but not across.

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

For mirroring to be optimal, two conditions have to hold. Profits have to be higher if each agent informs the other agents in his own module, and no one else, than if he engages in no communication at all. Moreover, conditional on each agent informing the other agents in his own module, profits have to be higher if the agent refrains from any further, across-module communication. To determine when these conditions are met, we can draw on the results in the previous section. For the first condition, we know from (9) that if agent 1 shares his state with the other agents in his module, but no one else, his expected revenue is given by

$$R_1\left(\boldsymbol{C}_1\left(\ell\right)\right) = a_1^2\sigma_1^2\left(1 + \frac{p_1^2\left(n_1 - 1\right)}{\left(1 + p_1\right)\left(1 - \left(n_1 - 1\right)p_1\right)}\right),$$

where the first term in brackets is the value of all walks from node 1 back to itself that only go through node 1 and the second is the value of all walks from node 1 back to itself that go through at least one other node in module $\mathcal{M}_1$ but does not go through any nodes in other modules. These latter walks are the ones that are created if agent 1 informs the other agents in his own module but no one else. The next lemma then follows.

LEMMA 5. *Profits are higher if each agent informs the other agents in his own module, and no one else, than if he engages in no communication at all if and only if*

$$\frac{a_i^2\sigma_i^2}{n_i - 1}\frac{p_i^2\left(n_i - 1\right)}{\left(1 + p_i\right)\left(1 - \left(n_i - 1\right)p_i\right)} \geq \gamma \text{ for all } i \in \mathcal{N}. \tag{11}$$

The term on the left-hand side–the per node marginal revenue that is generated by agent i informing the other agents in his module–is increasing in n_i and p_i. For mirroring to be optimal, each module, therefore, has to have enough members, and coordination among them has to be sufficiently important.

At the same time, each module cannot have too many members, and coordination among them cannot be too important. If they did, it would be optimal to inform the agents in those modules about states in other modules. This would violate the second condition for mirroring to be optimal, the absence of any communication across modules. Such across-module communication can never

be optimal when the degree of coupling is sufficiently low. The more members modules have, though, and the more important it is to coordinate among them, the lower the degree of coupling needs to be for the absence of across-module communication to be optimal.

LEMMA 6. *Suppose each agent $i \in \mathcal{N}$ informs the other agents in his module $\mathcal{M}_{m(i)}$ about his state. There exists a threshold degree of coupling $\bar{t}_i > 0$ such that profits are higher if agent i refrains from informing agents in other modules if and only if $t \leq \bar{t}_i$. The threshold $\bar{t}_i$ is decreasing in the number of modules M and in the module characteristics $n_{m'}$ and $p_{m'}$ for all $\mathcal{M}_{m'} \in \mathcal{M} \backslash \mathcal{M}_{m(i)}$.*

The conditions under which mirroring is optimal then follow directly.

PROPOSITION 3. *Mirroring is optimal if and only if* (11) *holds for all $i \in \mathcal{N}$ and $t \leq \min_{i \in \mathcal{N}} \bar{t}_i$.*

The proposition shows that for the mirroring hypothesis to hold, the conditions have to be just right. There cannot be too many modules and none of the modules can consist of too many decisions or require too much coordination, for else some agents should share their information beyond their own modules. At the same time, there also cannot be modules that consist of too few members or require too little coordination, for else the agents in such modules should not even share their information with each other. Mirroring, in other words, is associated with moderation. Modular production favors modular organization only if there is a limited number of modules, each of which consists of an intermediate number of decisions that require an intermediate degree of coordination.

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

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

Partial monitoring can be optimal in our setting and Figure 4 provides an example. If it is, though, it has to take a particular form.

PROPOSITION 4. *When partial mirroring is optimal, the organization contains one cluster of modules and the modules that form the cluster are those $\mathcal{M}_m \in \mathcal{M}$ with the largest values of x_m.*

The result follows from the optimality of hierarchies we discussed in the previous section. If there were multiple clusters, hierarchies would not be nested, which cannot be optimal. Suppose,

for instance, that one cluster consists of modules $\mathcal{M}_1$ and $\mathcal{M}_2$ and another of modules $\mathcal{M}_3$ and $\mathcal{M}_4$. If it is optimal for an agent in $\mathcal{M}_1$ to report to agents in $\mathcal{M}_2$ but not to those in $\mathcal{M}_3$, then $\mathcal{M}_2$ has to be more cohesive than $\mathcal{M}_3$. But if $\mathcal{M}_2$ is more cohesive than $\mathcal{M}_3$, it cannot be optimal for an agent $\mathcal{M}_4$ to report to agents in $\mathcal{M}_3$ but not those in $\mathcal{M}_2$. In contrast, the existence of a single cluster is consistent with the optimal design of communication networks, provided it consists of the most cohesive modules. The optimal communication network then has a core-periphery structure in which the core consists of the most cohesive modules, whose agents all communicate with each other, and the periphery of the less cohesive ones, whose agents only communicate with others in the same module.