Communication Networks: Knowledge and Decisions

Antoni Calvó-Armengol and Joan de Martí *

That communication pervades the everyday life of organizations is a statement of the obvious. However, how this widespread communication exactly affects organization's performance is less clear. Another truism about organizations is that their members are seldom fully knowledgeable of the exact state of the world under which they operate. To overcome this information uncertainty burden, organizations tend to devote a substantial amount of resources to retrieve information about their environment.

Individual experimentation is a natural means to obtain superior information. When the nature and the consequences of the prevailing uncertainty is common to all organization members, a natural substitute to replications of individual experiments is to have agents communicating their private information among them. Communication is then a means to save on experimentation costs when these display diminishing returns.

In another vein, shared information fosters coordination among individual agents, which is beneficial to the organization as a whole when coordinated actions lead to higher benefits for everyone. Altogether, when each organization member individually ascertains that both uncertainty and coordination are the main driving forces for his activity, communication generates complementarities in all dimensions of the decision process.

Individual experimentation generates private information dispersed inside the organization. How this dispersed individual information is aggregated at the level of each agent depends on the fine details of the communication process available to the organization. Given an arbitrary communication process, individual optimal actions depend on the information structure this one generates inside the organization, since how information is pooled together shapes the beliefs each member has about the information accumulated by other agents. We analyze how communication moulds actions, and illustrate this link for a broad class of networked communication processes.

The networked class of decentralized communication devices we consider is the following. Agents connected through a given network, and only them, communicate in pairs and for a fixed number of rounds. At each round, agents average the stream of signals previously received from their network contacts, and communicate this average signal back to them. Averaging reduces the volume of information prior to its retransmission, and thus saves on information-processing costs. However, from round two onwards, this simple heuristic that treats signals in the stream as *de facto* mutually independent, fails to adjust properly for redundant information from a common third-party. In fact, Peter DeMarzo, Dimitri Vayanos and Jeffrey Zwiebel 2003 show that the converging beliefs dynamics for this simple rule over-weight the private information of more "central" agents in the communication network. For this reason, asymptotic beliefs are not correct when the underlying network is irregular.

Here, we analyze how this decentralized information-sharing scheme shapes individual and organization decisions and outcomes when the network geometry varies and for arbitrary communication rounds.

We combine the communication device with a model in which each agents' utility includes a common decision problem, and a coordination premium that aligns the individual interests of all the organization members. In this setup, we characterize the unique Bayes-Nash equilibrium for general communication processes. The equilibrium strategies are linear and depend on the knowledge index of each agent, computed for the information structure induced by the prevailing communication process. The knowledge index, which we introduce, measures the (arbitrary) high-order beliefs each agent entertains about every other agents' information, and thus reflects the communication possibilities available to everyone. This index is formally reminiscent of standard centrality measures in sociology, but is computed with an information correlation matrix rather than with socio-metric data.

Aggregate *ex ante* equilibrium payoffs, which we express directly in terms of knowledge indexes, allow for a welfare assessment of general communication processes (compared to the first-best world without information uncertainty). For the class of decentralized information-sharing schemes described above, we document the role of the geometry of the network of communication channels, and that of the number of available communication rounds.

The payoffs induced by an information structure, however, only give a partial assessment of its performance. A notion of cost associated to every information structure is also required to perform a proper evaluation and cross comparisons. For instance, Kenneth Arrow 1985 suggests the entropy (an index of the distance between a probability distribution and certainty) as a possible cost measure. This, however, remains an abstract cost notion.

Our class of communication processes, instead, is procedurally described by its two main ingredients, the underlying network of open communication channels and the number of active communication rounds. We can thus identify two different and complementary cost notions for each communication process in this class, directly related to its specificities, namely, the installed communication capacity (the links), and the exact length of communication (the rounds). As a matter of fact, the experimental psychology literature has long ago documented the crucial role of communication time and communication pattern for information aggregation purposes, as reported in the seminal work by Alex Bavelas 1950.

Our equilibrium characterization with knowledge indexes, and the subsequent expression for aggregate payoffs, is fairly general and can potentially accomodate every conceivable communication process beyond the relatively broad class that we analyze here. In particular, we have invoked information-processing costs to justify averaging away the incoming streams of signals prior to their retransmission. Instead, one could imagine a more general transformation of streams of signals with its associated time or cognitive costs, and use our results to readily draw conclusions about its overall performance.

I. Communication and forecasts

The underlying state of the world θ is normally distributed with mean zero and variance ϕ_θ. There is a private signal $x_i = \theta + \varepsilon_i$ for each of the n players, conditionally independent. The noise ε_i is normal with mean zero and variance ϕ_ε. Players share their private signals according to some communication process $\mathbb{P}$, which delivers some output signals $y_i^{\mathbb{P}} = \theta + \varepsilon_i^{\mathbb{P}}$. We assume that the vector of output noises $\boldsymbol{\varepsilon}^{\mathbb{P}}$ is multi-normal with mean $\mathbf{0}$ and variance-covariance matrix $\phi_\varepsilon \boldsymbol{\Sigma}^{\mathbb{P}}$, with coefficients $\phi_\varepsilon \sigma_{ij}^{\mathbb{P}}$.

Based on his information, each of the n players chooses an action a_i, and incurs losses from a mismatch between this choice, the parameter θ and others' choices (averaged):

(1.1) $$u_i(\mathbf{a};\theta) = -(1-r)(a_i-\theta)^2 - r\frac{1}{n-1}\sum_{j\neq i}(a_i-a_j)^2.$$

The parameter r, between zero and one, weights both sources of quadratic loss, related to the accuracy in predicting the state of the world and in inferring others's choices. The distribution of output signals determines this ability to forecast and to draw inferences.

Let $\alpha = \phi_\varepsilon / \phi_\theta$. Denote by $\mathbb{E}_i^{\mathbb{P}}(\bullet)$ the conditional expectation operator by player i when his output signal is $y_i^{\mathbb{P}} = \hat{y}_i$. Standard algebra with normal distributions gives $\mathbb{E}_i^{\mathbb{P}}(\theta) = f_i^{\mathbb{P}}\hat{y}_i$ and $\mathbb{E}_i^{\mathbb{P}}(y_j^{\mathbb{P}}) = \omega_{ji}^{\mathbb{P}}\hat{y}_i$, where:

(1.2) $$f_i^{\mathbb{P}} = \frac{1}{1+\alpha\sigma_{ii}^{\mathbb{P}}} \text{ and } \omega_{ji}^{\mathbb{P}} = \frac{1+\alpha\sigma_{ij}^{\mathbb{P}}}{1+\alpha\sigma_{ii}^{\mathbb{P}}}, \text{ for all } i, j.$$

Let $\rho = r/(n-1)$. Given that payoffs are quadratic and concave in own action, the optimal action that maximizes expected payoffs conditional on own output signal (and thus solves first-order conditions) is linear in the forecast and in the inferences based on this output signal:

(1.3) $$a_i^{\mathbb{P}}\left(\hat{y}_i\right) = (1-r)\mathbb{E}_i^{\mathbb{P}}(\theta) + \rho\sum_{j\neq i}\mathbb{E}_i^{\mathbb{P}}\left(a_j\right).$$

Plugging back the expression for the optimal action $a_j^{\mathbb{P}}$ into the expression for the optimal action $a_i^{\mathbb{P}}$, and recursively iterating the process, leads to $a_i^{\mathbb{P}}$ as an infinite sum of arbitrary high-order iterated expectations $\mathbb{E}_i^{\mathbb{P}}\mathbb{E}_{i_2}^{\mathbb{P}}\cdots\mathbb{E}_{i_q}^{\mathbb{P}}(\theta)$ over the state of the world θ. Such iterated expectations involve the inferences players make about others' forecasts of the state of the world, themselves related to the inferences players make about each others' output signals. The linearity of the expectation operator readily implies the following:

(1.4) $$\mathbb{E}_i^{\mathbb{P}}\mathbb{E}_{i_2}^{\mathbb{P}}\cdots\mathbb{E}_{i_q}^{\mathbb{P}}(\theta) = f_{i_q}^{\mathbb{P}}\omega_{i_q i_{q-1}}^{\mathbb{P}}\cdots\omega_{i_2 i_1}^{\mathbb{P}}\hat{y}_{i_1}.$$

Arbitrary high-order iterated expectations depend both on the number q of such iterations, and on all the possible ordered identities $i_1, i_2, \ldots, i_q$ of the players along the chain of pair-wise inferences entering these iterated expectations (expectations about whom expectations, about whom expectations ...). These high-order beliefs reflect the ability of each player to infer the information held by the other players after communication, and thus depend on the details of the communication process $\mathbb{P}$, and the corresponding correlation in output signals in $\phi_\varepsilon \mathbf{\Sigma}^{\mathbb{P}}$.

In general, we cannot invoke nor construct some average belief operator to compute high-order beliefs, neither expect symmetric behavior from the part of the players. Rather, high-order iterated expectations are an-isotropic, and change with the particular ordered chain of pair-wise inferences. This an-isotropy very likely sustains asymmetric choices across players.

II. Knowledge index and decisions

We now characterize the Bayes-Nash equilibria of the game when players' information is given by the output signals after communication.

Notice that high-order beliefs by player i in (1.4) are a linear function of his output signal realizations $\hat{y}_i$. This suggests an equilibrium strategy linear in own information. In another vein, high-order beliefs of order q are obtained by nesting q linear best-responses (1.3), and are thus discounted by a factor ρ^q. This, in turn, suggests that the coefficient in the (candidate) linear equilibrium may involve high-order beliefs discounted by an appropriate ρ^q.

Let $\mathbf{\Omega}^{\mathbb{P}}$ be a zero diagonal matrix with out-of-diagonal terms $\omega_{ij}^{\mathbb{P}}$. Then, the coefficients of the qth power of this matrix are the $\omega_{i_q i_{q-1}}^{\mathbb{P}} \cdots \omega_{i_2 i_1}^{\mathbb{P}}$ that enter the calculation of order q beliefs in (1.4). The following vector keeps track of the discounted high-order beliefs of arbitrary order that arise when linear best-responses (1.3) are nested into each other *ad infinitum* to compute the equilibrium fixed point ($\mathbf{1}$ is the vector of ones.):

$$\mathbf{1} + \rho \mathbf{\Omega}^{\mathbb{P}} \cdot \mathbf{1} + \rho^2 \left(\mathbf{\Omega}^{\mathbb{P}} \right)^2 \cdot \mathbf{1} + \rho^3 \left(\mathbf{\Omega}^{\mathbb{P}} \right)^3 \cdot \mathbf{1} + \cdots = \left(\mathbf{I} - \rho \mathbf{\Omega}^{\mathbb{P}} \right)^{-1} \cdot \mathbf{1}.$$

This is a non-negative vector, and each coordinate reaches its upper bound $(1-r)^{-1}$ when input signals are fully informative ($\phi_\varepsilon = 0 = \alpha$).

Definition. *The knowledge index for the communication* $\mathbb{P}$ *is* $\mathbf{k}^{\mathbb{P}} = (1-r)\left(\mathbf{I} - \rho \mathbf{\Omega}^{\mathbb{P}} \right)^{-1} \cdot \mathbf{1}$.

This is a normalized vector whose coordinates between zero and one give the discounted high-order beliefs of individual players for the current communication process among them.

Theorem. *The unique Bayes-Nash equilibrium for the communication process* $\mathbb{P}$ *has linear strategies* $a_i^{\mathbb{P}}\left(\hat{y}_i \right) = k_i^{\mathbb{P}} \mathbb{E}_i^{\mathbb{P}}(\theta) = k_i^{\mathbb{P}} f_i^{\mathbb{P}} \hat{y}_i$. *The equilibrium aggregate ex ante payoffs are:*

$$U^{\mathbb{P}} = (1-r)\phi_\theta \left[\sum_{i=1}^{n} f_i^{\mathbb{P}} \left(k_i^{\mathbb{P}}\right)^2 - n \right] \leq 0.$$

The uniqueness and linearity result follow from a central theorem by Roy Radner 1962 on teams, and the fact that our quadratic game payoffs admit a potential that represents common (team) interests for all players. The particular closed-form for the equilibrium strategy, that involves explicitly the knowledge index, exploits the intimate connection between quadratic games and centrality indexes in sociology established by Coralio Ballester, Antoni Calvó-Armengol and Yves Zenou 2006.

III. A class of networked communication processes

So far, the communication process is characterized by the distribution over its output signals. We now describe a particular instance of a communication process for which we can explicitly compute this distribution over output signals.

Let each player perform separately a different experiment. We get a particular realization of private signals $\hat{x}_i = \theta + \hat{\varepsilon}_i, i = 1,...,n$. If further experimentations are costly, information pooling is then a natural outcome –given the common interests in (1.1). Players exchange their private signals in pairs, and some given communication network g specifies who talks with whom. More precisely, $g_{ij} = g_{ji} = 1$ if i and j talk to each other, and $g_{ij} = g_{ji} = 0$, otherwise. Of course, $g_{ii} = 1$, while $g_i = g_{i1} + \cdots + g_{in}$ is the total number of interlocutors to player i, including oneself.

When this first communication round is over, each individual gets a stream of private signal realizations, one for each network interlocutor. Given normal distributions, it turns out that the average stream signal is a sufficient statistics to assess the value of the fundamental:

$$\mathbb{E}\left(\theta \middle| \left\{\hat{x}_j\right\}_{g_{ij}=1}\right) = \mathbb{E}\left(\theta \middle| \hat{x}_i^1\right), \text{ where } \hat{x}_i^1 = \frac{1}{g_i}\sum_{j=1}^{n} g_{ij}\hat{x}_j.$$

Denote by $\mathbf{G}=[g_{ij}]$ the adjacency matrix of communication links, and by $\mathbf{C}(g)=[c_{ij}(g)]=[g_{ij}/g_i]$ its row normalization. A compact notation for the average stream signal is $\hat{\mathbf{x}}^1=\mathbf{C}(g)\hat{\mathbf{x}}$.

At equilibrium, players forecast the value of θ, and $\hat{x}_i^1$ is a sufficient statistics for that matter. They also anticipate adequately each others' actions, which involves a forecast of each others' private information, forecasts about forecasts, etc. On top of easing retransmission, reducing the information contained in a stream of signals into a simple average renders these cross forecasts easier to compute. However, this inevitably comes with a loss in information compared to the straight cross forecasts based on the full signal streams.

When players communicate repeatedly with their direct interlocutors in the network g, and average the incoming stream of signals before re-transmitting it back to them, the resulting output signals after t completed rounds of talks are $\hat{\mathbf{x}}^t=\mathbf{C}^t(g)\hat{\mathbf{x}}$. We denote by $\mathbb{P}^t(g)$ this communication process.

A useful analogy with $\mathbb{P}^t(g)$ is the following. Let players pick a message recipient uniformly among all their possible interlocutors in g. We get a random walk on the network, where transmission from i to j occur with probability $c_{ij}(g)=g_{ij}/g_i$. More generally, the row-normalized matrix $\mathbf{C}^t(g)=[c_{ij}^{[t]}(g)]$ gives the probabilities of a random walk of length t between every pair i and j. The intersections and partial overlaps across of all the $t-$random walks emanating from i and j then determine the correlation between the output signals x_i^t and x_j^t that these players get after t completed rounds of conversations:

(3.1) $$\phi_\varepsilon \sigma_{ij}^{\mathbb{P}^t(g)}=\sum_{k=1}^{n} c_{ik}^{[t]}(g)\, c_{jk}^{[t]}(g), \text{ for all } i, j.$$

The pattern of correlations is shaped by the geometry of the communication network. For instance, with one communication round, the correlation between two players' output signals is proportional to the fraction of common interlocutors they share among their communication partners.

Plugging (3.1) into (1.2) gives all the necessary ingredients to compute the Bayes-Nash equilibrium of $\mathbb{P}^t(g)$ for all g and t.

We compare the corresponding equilibrium payoffs for two different network geometries and for various communication rounds. To get a sense of the cost for the installed communication capacity, we compare two networks on $n=4$ agents with identical total number of links, the *kite* and the *wheel*.

[Insert Figure 1]

A simple linear expansion of the knowledge index implies that equilibrium payoffs are monotone with $\sum_{i=1}^{n}\sum_{j=1}^{n}\sigma_{ij}^{\mathbb{P}^t(g)}/\sigma_{ii}^{\mathbb{P}^t(g)}$ (for small r and high α), a measure of network irregularity. With one communication round, this index is worth 38/3 for the *kite* and 36/3 for the *wheel*. Payoffs are thus higher with $\mathbb{P}^{t=1}(kite)$ than with $\mathbb{P}^{t=1}(wheel)$ (for adequate α, r).

However, these payoffs ranking is reversed with enough communication rounds. The reason is the following. A message initially at i is at j at round t with probability $c_{ij}^{[t]}(g)$. Writing $\mathbf{C}^{t+1}(g)=\mathbf{C}(g)\mathbf{C}^t(g)$, one can view $\mathbf{C}(g)$ as the Markov transition matrix for the row probability vectors $\left(c_{i1}^{[t]}(g),\ldots,c_{in}^{[t]}(g)\right)$ of the row-normalized matrix $\mathbf{C}^t(g)$. We thus have $\lim_{t\to\infty} c_{ij}^{[t]}(g)=c_j^{\infty}(g)$, where $\mathbf{c}^{\infty}(g)$ is the unique invariant distribution of the irreducible and aperiodic Markov process with transitions $\mathbf{C}(g)$. In turn, the fact that all row vectors of $\mathbf{C}^{\infty}(g)$ are identical implies that long-run beliefs for $\mathbb{P}^{\infty}(g)$ are common to all players, that is:

$$\hat{x}_i^{\infty} = \hat{x}^{\infty} = c_1^{\infty}(g)\hat{x}_1 + \cdots + c_n^{\infty}(g)\hat{x}_n, \text{ for all } i,$$

a weighted sum of private input signals.

We compute the weights. With an un-directed network, $g_i c_{ij}^{[t]}(g) = c_{ji}^{[t]}(g) g_j$ from simple algebra, from which $g_i c_j^{\infty}(g) = c_i^{\infty}(g) g_j$ at the limit, and thus $c_i^{\infty}(g) = g_i / (g_1 + \cdots + g_n)$. Because averages of incoming signal streams at each communication round do not discount properly for redundant information from common sources, better connected players in the communication network end up credited with a higher weight in the emergent long-run consensual beliefs.

When all players share the same beliefs, the knowledge index hits its upper bound of one. Aggregate long-run payoffs then take the following simple form:

$$U^{\mathbb{P}^{\infty}(g)} = n(1-r)\phi_{\theta}\left[\frac{1}{1+\alpha\frac{g_1^2+\cdots+g_n^2}{(g_1+\cdots+g_n)^2}}-1\right]$$

Fix the total supply of links. Then, these payoffs are maximal when $g_1^2 + \cdots + g_n^2$ is minimal, namely, on a regular network. Aggregate payoffs are thus higher with $\mathbb{P}^{t=\infty}(wheel)$ than with $\mathbb{P}^{t=\infty}(kite)$.

[Insert Figure 2]

In a regular network, $c_i^{\infty}(g) = 1/n$ for all players. All private signals are weighted evenly, and there is no scope for biased social influence by anyone on anyone else. In an irregular network, instead, weights vary across players with their network connectivity. These unequal weights do not wash out any possible distortion in the independent private input signals, but rather amplify it, thus resulting in an un-ambiguous welfare loss. DeMarzo, Vayanos and Zwiebel 2003 already show that "un-balanced" (irregular) network geometries yield to incorrect long-run beliefs. Here, we gain an exact welfare assessment of this bias.

Arrow, Kenneth. 1985. "Informational Structure of the Firm." *American Economic Review, Papers and Proceedings* 75(2): 303-307.

Ballester, Coralio, Antoni Calvó-Armengol and Yves Zenou. 2006. "Who's Who in Networks. Wanted: the Key Player." *Econometrica* 75: 1403-1418.

Bavelas, Alex. 1950. "Communication Patterns in Task-Oriented Groups." *Journal of the Acoustical Society of America* 22:725-730.

Calvó-Armengol, Antoni and Joan de Martí. 2006. "On Optimal Communication Networks." http://selene.uab.es/acalvo

DeMarzo, Peter, Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics* 118: 909-968.

Morris, Stephen and Hyun Song Shin. 2006. "Optimal Communication." http://www.princeton.edu/~smorris/

Radner, Roy. 1962. "Team Decision Problems." *Annals of Mathematical Statistics* 33: 857-888.

Sobel, Joel. 2006. "Information Aggregation and Group Decisions." http://weber.ucsd.edu/~jsobel/

* Antoni Calvó-Armengol (corresponding author): ICREA, Universitat Autònoma de Barcelona, and CEPR. Department of Economics, Edifici B, 08193 Bellaterra (Barcelona), Spain. Email: antoni.calvo@uab.es ; Joan de Martí: Department of Economics, Universitat Autònoma de Barcelona, Edifici B, 08193 Bellaterra (Barcelona), Spain. Email: joan.demarti@idea.uab.es. Sébastien Bervoets provided excellent research assistance. Financial support from the Fundación BBVA, the Spanish Ministry of Education and FEDER through grant SEJ2005-01481ECON is gratefully acknowledged.

Figure 1.

Two networks with four links and four players each, the kite and the wheel.

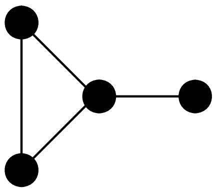

Figure 2.

Ex ante equilibrium aggregate payoffs from round 1 to round 5 for the kite and for the wheel

(computed for $r = 0.5$ and $\alpha = 5$).