

Diversity and Communication in Teams: Improving Problem-Solving or Creating Confusion?

M. Laura Frigotto · Alessandro Rossi

Published online: 7 May 2011
© Springer Science+Business Media B.V. 2011

Abstract Despite the rich and interdisciplinary debate on the role of diversity and communication in group problem-solving, as well as recognition of the interactions between the two topics, they have rarely been treated as a joint research issue. In this paper, we develop a computational approach aimed at modeling problem-solving agents and assess the influence of various levels of diversity and communication in teams on agents' performance in problem-solving. By communication, we intend a conversation on the persuasiveness of the features characterizing problem-setting. By diversity, we mean differences in how agents build problem representations which allow them to access various solutions. We deploy the concept of diversity along two dimensions: knowledge amplitude, which accounts for the level of available knowledge allowing access to poorer or richer problem representations (compared with complete problem representations), and knowledge variety, which pertains to the differences in the constituents of agents' representations. We define performance as the frequency with which diverse agents choose the same alternative representation of an agent displaying complete representations of the problem. Our results indicate that communication is more effective when agents elaborate from relatively richer problem representations, as this provides a basis for integrating the variously diverse beliefs of their teammates. Conversely, poorer diverse representations may lead to worse performance when knowledge variety also applies. Lastly, we show that the influence of communication is not monotonically positive, as increasing communication intensity performance may worsen at any level of knowledge availability and knowledge variety.

M. L. Frigotto · A. Rossi (✉)
Department of Management and Computer Science, University of Trento, Trento, Italy
e-mail: alessandro.rossi@unitn.it

M. L. Frigotto
e-mail: marialaura.frigotto@unitn.it

Keywords Problem-solving · Diversity · Heterogeneous agents · Communication · Constraint satisfaction · Neural networks · Causality

1 Introduction

An increasingly more diverse society poses a challenge to how work, education and interpersonal relationships are conceived and organized. It is no longer possible to think of communities, organizations and teams as homogeneous collections of typical individuals. This kind of pervasive and endogenous heterogeneity brings new interest to questions such as what diversity means and how to deal with it. Conversely, the organization of the workforce has moved toward less hierarchical structures based on direct communication and interaction of peers, such as interactive groups and problem-solving teams (Ilgen 1999; Jackson et al. 1992). In this work, we examine to what extent and in what kinds of conditions interactions between diverse individuals enhance human performance in problem-solving activities. We analyze a problem-solving task performed by a set of diverse agents with bounded abilities in collecting the relevant features which characterize problem-setting and compare their individual performances with that of a team with increasing levels of communication.

In the previous literature, the issue of effectiveness has been explored in teams of problem-solvers along the dimensions of communication and diversity taken separately. Some studies (Hutchins 1995; Marchiori and Warglien 2005) have explored to what extent communication can lead to better and more rapid solutions. Other contributions have addressed the role of diversity in problem-solving (e.g., Hong and Page 2001, 2004), but have neglected considering communication as a moderating variable or have even excluded it from analysis. An extreme exemplar is provided by the model of Hong and Page (2001, 2004), in which diverse teams are modeled as a collection of agents working sequentially on a common task, without introducing any form of communication.

In most of the literature, the discourse on communication is less focused on measuring performance, being more concentrated on understanding how individuals form groups, who they decide to team with and why (e.g. in the experimental literature: Kocher et al. 2006; Huck and Rey-Biel 2006; Page et al. 2005; Geard and Bullock 2008; Keser and Montmarquette 2007; Ahn et al. 2008), which role they assume in the relationship (Morgeson and Hofmann 1999; Kozlowski and Klein 2000; Schneider et al. 2000), how they manage and maintain the group, and how the team member is changed (in terms of behavior or attitude) by social interactions (Trenholm 1986). Nevertheless, some contributions have indicated that communication displays a non-monotonic relation with performance. In fact, over a certain level, communication strength influences the ability of the individual to think correctly. In some models, this effect has been represented as high pressure to conform to whatever the shared outcome may be—as suggested, for instance, by the “Credulous Theorem” of Marchiori and Warglien (2005). Similarly, in other models, high levels of communication result in the emergence of confusion in agents’ judgments, up to the point at which they are unable to select one out of many alternatives (Frigotto and Rossi 2006).

In a similar way, the contribution of diversity to collective problem-solving turns out to be controversial in the existing literature. In the decision-making approach (see [Williams and O'Reilly 1998](#) or [Mannix and Neale 2005](#), for a review), diverse agents are meant to offer access to different networks, information and expertise, i.e., different perspectives on problems, which increase problem-solving success ([Bantel and Jackson 1989](#); [Ancona and Caldwell 1992](#); [Winquist and Larson 1998](#); [Wittenbaum and Stasser 1996](#); [Carpenter 2002](#); [Pitcher and Smith 2001](#); [Kilduff et al. 2000](#); [Jackson et al. 2003](#)). Instead, the similarity-attraction paradigm indicates the limitations of diversity: people face obstacles in sharing thoughts and negotiating meanings if they do not have some background knowledge in common. As a result, irreconcilability or disregard of each other's ideas are more likely to appear and conflict is more likely to emerge ([Chatman and Flynn 2001](#); [Jehn et al. 1999](#); [Pelled et al. 1999](#); [Watson et al. 1993](#); [Webber and Donahue 2001](#)). Accordingly, problem-solving in teams reveals outcomes which are worse than individual ones, meaning that the exchange of ideas among very different people confuses thoughts instead of illuminating minds.

As several reviews have claimed ([Mannix and Neale 2005](#); [Williams and O'Reilly 1998](#); [Milliken and Martins 1996](#); [Guzzo and Dickson 1996](#)), such opposite conclusions on the role of diversity on performance derive from the differing definitions used (surface-level differences such as race/ethnicity, gender or age vs. underlying differences such as professional background, education or personality), research methods adopted (laboratory settings vs. real working groups) and the nature of the problems addressed (innovative/creative vs. conventional). However, contrasting findings stemming from the various paradigms are evidence of a unifying phenomenon which requires less fragmented analysis. We believe that divergent results deriving from research on diversity and communication, taken separately, may be partially clarified by joint assessment of their interplay. This attempt is probably still weak, when viewed in relation to the complexity of group communication and the multifaceted nature of diversity, which can only partially be captured by simplifications of broader categories. As a matter of fact, scholars from several disciplines (psychology, economics, sociology, anthropology, communication and education research and the organizational field) have provided a plurality of research efforts aimed at unveiling the origins, meaning and dynamics of diversity and communication at various levels of analysis and epistemological perspective. However, this plethora of contributions has not produced a clear picture of these phenomena, and there is a lack of unity of meaning and of metric ([Mannix and Neale 2005](#); [Williams and O'Reilly 1998](#); [Trenholm 1986](#); [Deetz 2001](#)) which results in the conceptual confusion typical of "overburdened concepts" as such ([Trenholm 1986](#)). In this context, we believe that joint stylization of communication and diversity may help to clarify some basic dynamics among variables which may provide intuitions for multidimensional analyses of diverse groups facilitated by other methodological tools. We believe that our understanding would benefit by knowing to what extent communication effects depend on agents' diversity and, vice versa, whether the role of diversity is affected by the strength of communication: this paper may be regarded as an attempt to address this gap in collective problem-solving. Our research question is twofold. On one hand, we ask to what extent diversity supports collective problem-solving, allowing effective communication among agents, or

whether it only adds confusion and “noise” to the interpretation of the problem-setting. On the other hand, we are interested in understanding whether more communication can usefully support lower levels of diversity among agents.

We address these questions through a computational model, deploying teams of diverse agents communicating with each other.

The rest of the paper is organized into six sections. Section 2 positions this contribution with respect to the existing models in the literature. Section 3 reviews some models of teams of agents addressing the problem of diversity and communication. Section 4 presents our model, and Sect. 5 displays the structure of our simulations. Section 6 is devoted to presentation of the results, which are finally discussed in Sect. 7.

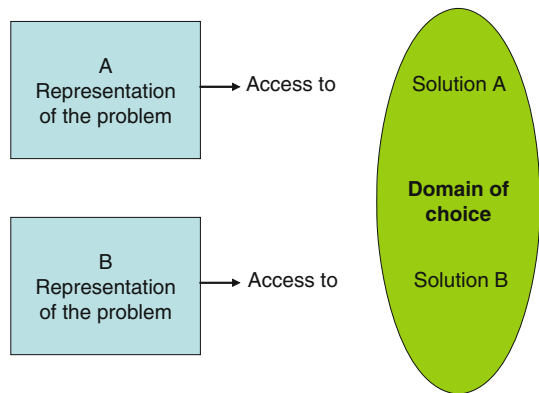
2 The Conceptual Model

Our model has three main novelties with respect to existing models in the literature. First, we model diversity in terms of differences in how agents interpret problems. The literature on diversity has referred to attributes such as functional background, education, training, experience or personality, which has proved to be the source of the “value-in-diversity hypothesis” (Williams and O’Reilly 1998). Conversely, the modeling tradition has represented this issue indirectly in terms of information exposure or information processing ability, and has compressed the two dimensions of information and knowledge by considering the first as a proxy and an effect of the second. Instead, in our model, we distinguish two levels: information, and knowledge, which allows us to interpret and benefit from information. In this way, we make a distinction between agents who have the same access to information but who cannot interpret it.

Second, we provide a measure of diversity, which allows us to assess the influence on performance at various levels of heterogeneity. Hong and Page (2001) noted a shift between models in which problem representations were fixed among agents, to a model in which they varied. However, they chose to model agents’ perspective on the problem through a binary string of a certain length, which does not allow us to assess the distance between two diverse agents having different binary strings in terms of representations. We address this limitation by modeling agents’ diversity in terms of two different dimensions, both referring to agents’ background knowledge, and not simply to information availability. We assume as our reference point the knowledge required for the complete representation of a problem, and model diversity in terms of different knowledge bases available to agents. We measure *knowledge amplitude* in terms of the number of causal links which compose problem representation, and *knowledge variety* in terms of differences in causal links. We then explore how agents’ interactions and problem-solving outcomes vary along these two features of diversity.

Third, our model belongs to a peculiar class of connectionist networks, i.e., constraint satisfaction (CS) networks, originally proposed by Rumelhart et al. (1986). Other contributions on team diversity have built upon this approach, assuming problem-solving as a pattern-matching activity (Hutchins 1995; Marchiori and Warglien 2005). The literature on diversity has shown its main value when faced with creative problems (Hoffman 1959; Hoffman and Maier 1961; Triandis et al. 1965; Nemeth 1986), which can hardly be reconciled with pattern recognition. Nevertheless,

Fig. 1 Problem representations and solutions



we still support the use of CS models, but suggest abandoning the pattern-matching perspective, thus extending the class of cognitive phenomena which can be represented through CS networks. In our model, agents do not recall a situation pattern when confronted with environmental stimuli, with a readily available solution attached to it; they engage in a search to build a causal explanation of the problem, by means of which they can understand how to act on it.

We assume that agents, viewed as problem solvers, struggle for building convincing and internally coherent representations of problems, based on a satisfactory understanding of the nature and relationships of the constituents of the problem setting. This means that they try to give sense to the scraps of evidence they consider important for the issue at stake, to give an explanation to the causes of their manifestation in order to understand how to intervene on the situation. As sketched in Fig. 1, we hypothesize that alternative representations give access to alternative solutions (Simon 1991), and that an assessment over the internal coherence of the representation makes them decide for the more convincing representation as well as for a solution to the problem setting.

In philosophy, the logic of explanation refers to the literature on explanatory coherence and to the substantial contribution of Thagard (1989, 1992b, 2000) which not only developed the concept theoretically but also provided numerous applications, for instance, in the domains of scientific revolutions (1989), scientific discoveries (1998a), medical discoveries (1998b), adversarial problem-solving (1992a) and juror's decisions (2004). The psychological literature has acknowledged the importance of explanation as a primitive source of knowledge, for example in children (Keil 2006). A similar function of explanation is reported in the organizational literature in the so-called High Reliability Organizations (HRO), i.e. organizations working in extremely hazardous environments, such as military organizations, nuclear power companies or aviation safety organizations. They are accustomed to building narratives and exploring explanations of their everyday experience, in order to expand their knowledge and preparation for the most risk-laden events which they have not (yet) experienced (March et al. 1991; Morris and Moore 2000; Weick and Sutcliffe 2001). Organization scholars have also highlighted the fact that explanations are important in the past

perspective, in order to explain and rationalize action, as a form of retrospective rationality (March 1975, 1994), but also in a forward-looking perspective, as a means for formulating a diagnosis (Patriotta 2003).

Before proceeding further, we believe it is worth mentioning several features which are not included in our model. First and foremost, our model is not one of search for solutions. We do not treat search problems, since our purpose is to focus on how diverse solutions can circulate among agents in such a way that each can benefit from another's ideas and knowledge, and eventually take a wiser decision. Our agents have already searched for their solutions, and the way they did so is not considered within the model. Second, our agents are also bounded rational in the sense that they can assess only two alternatives at a time. We restrict our analysis to what they choose between two representations, in view of the way in which they build them.

3 Modeling Agents' Diversity, Information Availability and Communication

There is a rich interdisciplinary debate on the role of diversity within teams of agents. A large proportion of contributions has focused on domain-specific models and empirical studies, targeted at better understanding the phenomenon within a particular field. In this section, we limit our analysis to a more general class of models of diversity in collective problem-solving, which have been developed at a more abstract level, without close reference to a specific field of application. Nevertheless, in many cases, the implications for the various domains of problem-solving can easily be derived.

In this section, we consider four models of diversity in group problem-solving. Recalling the distinction made by Newell and Simon (1972) between problem-solving and problem-setting. We discuss contributions according to their source of diversity. Specifically, we distinguish diversity deriving from differences in problem-solving, problem-setting and information availability among agents.

In the models of Hong and Page (2001, 2004), bounded rational agents search in a landscape of solutions for the configuration displaying the maximum fitness value. Diversity is introduced by modeling agents with heterogeneous problem-solving abilities in a twofold way: through (i) different problem-setting ("perspectives") and/or (ii) different problem-solving strategies ("heuristics"). The above authors do not consider such questions as asymmetric and imperfect information, as the available landscape for searching is the same for every agent. Rather, they conceive diversity as emerging from differences in knowledge at two distinct levels: (i) in the 2004 paper, and both (i) and (ii) in the 2001 paper.

In their 2004 paper, the authors claim that the more diverse agents are—with respect to heuristics—the better the group performance, due to the larger landscape that can be explored. Results show that individual problem-solvers outperform the group of the best performing agents. Instead, in their 2001 paper, Hong and Page examined diversity in terms of pairs of perspectives and heuristics. In particular, diverse perspectives imply that agents translate their landscape into a problem space which is unique. As a result, the authors show that, *ceteris paribus*, diversity in perspectives enlarges the set of solutions considered during the search process. Note that in these models agents do not communicate with each other, but rather are considered as a set

of independent problem-solvers operating sequentially or simultaneously on the same problem. Teams are defined with reference to each agent's performance: groups result from evaluation of the average individual and totally autonomous performance of n best agents with regard to n randomly extracted agents, where diversity is assumed to derive from random selection.

Communication and information availability are central to the works of [Hutchins \(1995\)](#) and [Marchiori and Warglien \(2005\)](#); both contributions model teams of agents through a connectionist approach. More traditionally, [Hutchins \(1995\)](#) considers diversity as externally generated by imperfect information, originating either from poor problem-setting or poor information availability. Both conceive the inferential process and, more broadly, problem-setting, as deriving from the action of comparing environmental stimuli with the knowledge agents have collected and stored in the form of memorized patterns. Once pattern-matching has identified a specific case which has been encountered before, choice derives almost automatically as a consequence of stored solutions attached to that setting.

In [Marchiori and Warglien \(2005\)](#) first model various agents receive different (noisy) signals from the environment, so that they identify different pattern cases to be applied, misperceiving the true state of the world. As a result, diversity stems from information availability and the authors show to what extent communication can correct this kind of erroneous problem-setting. In a second and third series of simulations, the authors model diversity as an inner characteristic of the agents: patterns are distorted or incomplete. Diversity here originates from incomplete knowledge or distorted storage regarding the patterns of experienced cases. The authors show that, in these cases, communication corrects problems, although the corrective power follows a non-monotonic trend.

In general, when models explore diversity, they display the underlying assumption that communication among team members has no dysfunctional facets (with the significant exception of Marchiori and Warglien's "Credulous Theorem", mentioned earlier). However, the question of diversity in teams opens the door to considerations on how effective the communication of ideas can be among diverse agents, especially when diversity is considerable i.e. is not only related to differences in information access but also to underlying knowledge differences. It is worth mentioning here that, in most of these models, the outcome considered is not a single solution shared by all the team members; rather, agents interact and exchange ideas, all maintaining the autonomy of taking their own decision.

Agents with identical knowledge bases should communicate with each another easily. However, problem-solvers having differing perspectives on the problem may have trouble in understanding each other's solutions. In order to balance potential benefits and mishaps related to the introduction of communication in diverse agents, it is worth addressing the problem of how much knowledge agents need to have in common, in order to understand each other and, conversely, how much diversity is best. In other words, how much background knowledge, including language and vocabulary, do they need to have in order to benefit most from each other's specialization, without wasting time and energy in search of a common basis to which to refer? We address these questions in the following sections, in which we also provide our definition of diversity.

4 The General Model

4.1 Introductory Remarks

We model problem-solving in teams using a constraint satisfaction approach, stemming from the contribution of the PDP Group (Rumelhart et al. 1986) and build on the ECHO model proposed by Thagard (1992b).

Agents' perspectives of a problem are made up of a schemata of causality, in which they explain why some evidence appeared in order to figure out how to intervene. Explanations are built on the basis of agents' knowledge, and they connect evidence to knowledge. Information is meant to report about collected evidence. Formally, in the model, knowledge and evidence are represented by units, and causal explanations by links among knowledge and evidence units. For every problem, each agent builds two alternative explanations which compete against each other. Agents assess these perspectives of the problem on the internal consistency of the argument. The more internally consistent the alternative, the more convincing it is, and thus the more likely to be selected as the preferred one. Subsequently, they implement the choice/action associated with that preferred alternative.

Let us specify, formally and substantially, how our model differs with respect to more popular CS models assuming pattern-matching activities. According to these models, a series of patterns X_i (for $i = 1, \dots, n$) are stored in the network by careful selection of weights, and the aim of the relaxation process is to assess the ability of the model to recognize the correct pattern X_i when the signal coming from the environment is noisy (the initial activation is $X_i + E$). Conversely, in our model, each weight is assigned on the basis of the existence of an explanatory relationship which links two different units of the network, according to a rule (explained in the next subsection); like other constraint satisfaction models, such as that of Axelrod (1997), the aim of the relaxation process is to observe what solution the model displays.

4.2 Structure of the Model

4.2.1 Agents

Agents are represented by their schemata, which denote how they make sense of the world on the basis of their available data and knowledge.

Agents are modeled as a causal map of two theories or alternatives, each composed by a series of hypotheses causally explaining pieces of evidence. It is assumed that these competing theories have been compiled/processed according to agents' logical reasoning, knowledge, and access to evidence.

Formally, an agent is modeled as a constraint satisfaction network of n units, representing either pieces of available evidence (otherwise known as evidence units) or hypotheses, giving causal explanations for one or several pieces of evidence (or units of hypothesis, explanatory units or explainers). Special evidence units (which are only connected with evidence units) are also introduced into the model in order to fit evidence units to positive activation values (see below for details).

Explanatory units are grouped into two competing sets (theories, say A and B), representing alternative interpretations of problem-setting. Thus, the network may be imagined as a three-layer graph, in which the top layer represents explanatory units belonging to theory A, the middle layer evidence units, and the bottom layer collects the explanatory units belonging to theory B.

4.2.2 Activation of Units

The initial activation of units represents the agent's original beliefs regarding the units—that is, the agent's preliminary confidence about the environmental evidence and the various constituents of the theories. Unit activation, which may take on values in the $[-1, 1]$ interval, are updated overtime according to the relaxation rule (see below); the fixed point that is reached at the end of this process represents the agent's final belief regarding the units of the model. This steady-state may highlight the fact that the agent favors one theory over the other, if all explanatory units from one theory—say A—are positive, and vice versa for the units of theory B. If such a configuration does not occur, the model does not give clear indications in terms of choice of one theory over the other, suggesting a case in which the agent does not believe that the collected evidence and/or supporting hypotheses are conclusive.

4.2.3 Connections

Connections or weights in each agent's network w_{ij} are set in order to reflect the competitive or cooperative relationship existing between two units of the network (see [Thagard 1992b](#), for a full explanation). This procedure follows the rationale according to which agents try to build theories as series of arguments which support each other. For the sake of simplicity, we restrict our analysis to simple direct causality, so that explanatory links can only be set between units of evidence and explanatory units, and longer causal chains are not considered.

Positive connections between units of evidence and explanatory units represent direct causal relations (e.g. event x is causally explained by hypothesis y) and their intensity is coded through positive weights, so that higher weights correspond to higher causal relationships.

In order to introduce positive feedback between co-hypotheses, positive connections are also introduced between explanatory units (from the same theory) which jointly explain the same set of evidence units.

Lastly, competitive relationships are modeled by introducing negative connections between explanatory units belonging to opposing theories, which jointly explain the same units of evidence.

4.2.4 Procedure for Initialization of Activation

Let $\mathbf{s} = (s_1, \dots, s_i, \dots, s_n) = (a_1, \dots, a_k, b_1, \dots, b_l, e_1, \dots, e_m)$ be the vector of the activation of all the units in the network, where k and l represent the number of explanatory units belonging, respectively, to theories A and B, and m is the number of units of evidence (note that $n = k + l + m$). Let $\mathbf{f} = (f_1, \dots, f_i, \dots, f_n)$ (with

$f_i = 0$ for $i = 1, \dots, k+l$) also be the vector of the activation of the special evidence units.

Let W be a $n \times n$ null matrix. Define α as the excitatory default value for assigning positive connections among units and β as the inhibition default value for assigning negative connections among units. Then, weights w_{ij} (for $i = 1, \dots, n$, $j = 1, \dots, n$) are assigned according to the following steps:

- Step 1. Positive connection between an explanatory unit and a unit of evidence: for each unit of type e which is causally explained by one or more explanatory units of type a :
- i. let i corresponds to the position of unit e in s ;
 - ii. compute the number r of explanatory units of type a which explain s_i ;
 - iii. for each of the r explanatory units of type a which explain s_i ,
 - a. let j correspond to the position of unit a in s ;
 - b. set $w_{ij} = w_{ji} = \alpha/r$;
 - iv. repeat step 1 for theory B.
- Step 2. Positive connections between explanatory units which belong to the same theory: for each couple of units of type a which are co-hypotheses (they jointly explain one or more units of evidence):
- i. let i correspond to the position of one unit of type a in s and j to the position of the other unit in s ;
 - ii. set

$$w_{ij} = w_{ji} = \sum_{e=k+l+1}^n w_{ie} \cdot I;$$

where:

$$I = \begin{cases} 1 & \text{if } w_{ie} = w_{je} \\ 0 & \text{otherwise} \end{cases};$$

- iii. repeat step 2. for theory B.
- Step 3. Negative connections between explanatory units belonging to different theories: for each couple of units, one belonging to theory A and the other to theory B, which competitively explain one or more units of evidence:
- i. let i correspond to the position of the unit of type a and j to the position of the unit of type b in s ;
 - ii. set p as the number of units of evidence which are jointly explained by s_i and s_j ;
 - iii. set q as the overall number of co-hypotheses (of types a and b) which jointly explain the units of evidence at step ii;
 - iv. set $w_{ij} = w_{ji} = \beta p/(q/2)$.

Note that W is symmetric, $w_{ii} = 0$ for $i = 1, \dots, n$ and $w_{ij} = 0$ for $i = 1 + k + l, \dots, n$ and for $j = 1 + k + l, \dots, n$.

4.2.5 Relaxation Rule

Unit activation values are updated through a connectionist algorithm which is meant to increase the degree of coherence of the network, in the sense that it performs a gradient-descent path toward levels of activation of the units which best satisfy constraints (see [Rumelhart et al. 1986](#), for a formal treatment of analogous issues).

At each iteration, unit activation levels are synchronously updated according to the following rule ([Thagard 1992b, 2000](#); [Rumelhart et al. 1986](#)):

$$s_j(t+1) = (1-d)s_j(t) + \begin{cases} \text{net}_j(\max - s_j(t)) & \text{if } \text{net}_j > 0, \\ \text{net}_j(s_j(t) - \min) & \text{otherwise} \end{cases} \quad (1)$$

where $s_j(t)$ is the activation of unit j at time t , d is a decay parameter which, at each iteration, weakens the activation value of every unit. Min and max represent the lower and upper boundaries of the unit activation and are generally set, respectively, at -1 and 1 .

Lastly,

$$\text{net}_j = \sum_i w_{ij}s_i(t) + f_j \quad (2)$$

is the net input to unit j , computed as the sum of the activation of all the units weighted by connections w_{ij} linking each of these units with unit j . Note also that, in the case of evidence units, their net input also includes the value of the corresponding special evidence unit.

It is worth mentioning that formal treatment of this model is still incomplete. In particular, there is no proof of convergence of the system to a stationary state, nor of coherence maximization since, through relaxation, the system might settle on a local maximum. However, there is a considerable body of literature ([Thagard 1989, 1992a,b](#); [Nowak and Thagard 1992a,b](#); [Eliasmith and Thagard 1997](#)) that has shown convergence toward fixed points in finite time. In the simulations reported in the next section, we will employ a choice of parameters consistent with previous literature and study the issues of convergence and of multiplicity of local maxima (fixed points).

4.2.6 Communication

A group is a set of p agents modeled as a “network of agents’ networks” ([Hutchins 1995](#); [Marchiori and Warglien 2005](#)). Communication between two agents is modeled by linking each unit s_j of one agent with the corresponding unit s_j of the other agent. In this respect, communication is intended here as a parsimonious activity of exchange of beliefs, in which only activation of the units, and not the whole schemata, is shared. This means that, in our model, agents exchange beliefs on how important or how credible a hypothesis or a piece of evidence is, without telling each other how they constructed their causal relations or what their entire causal map looks like.

It is worth noting that our approach is far from being the only way of introducing communication between agents: [Hutchins \(1995\)](#) and [Marchiori and Warglien \(2005\)](#)

experimented various settings, varying symmetry and introducing hierarchic structures. Other ways of modeling communication among agents have been explored in other domains: for instance, [Thagard \(2000\)](#) studied how scientific consensus is reached by modeling agents who exchange verbal inputs, and [Thagard and Kroon \(2006\)](#) introduced emotional connections between agents.

In the present model, the vector of units is represented by the union of p agents' vectors of units (s), whereas the weight matrix contains both the individual weight matrices (arranged as $n \times n$ blocks along the main diagonal) and the communication matrices (arranged as $n \times n$ blocks outside the main diagonal). The strength of communication is modeled through the communication intensity parameter $\delta \geq 0$; note that, for $\delta = 0$, no communication occurs, and the model is reduced to a mere collection of independent agents. Also, as mentioned above, we assume the simplest form of communication: each agent communicates with all the others with the same strength (each communication matrix has $\alpha\delta$ over the main diagonal and 0 elsewhere). This model is still a constraint satisfaction network and we apply the same relaxation rule for modeling the individual case.

4.2.7 Diversity

The first source of diversity to be modeled is represented by differences in the number of causal explanations which agents might possess. The larger the set of explanatory links available to one agent, the larger that agent's knowledge amplitude.

By limiting the available explanatory links, we can model two different types of restrictions in agents' cognitive capabilities related to problem-solving. If one causal link is missing, it might be the case that the agent, albeit exposed to one evidence unit (already explained by one or more other explanatory units) and having already expressed a hypothetical unit (which already explains one or more evidence units) is not able to formulate correctly the existence of a causal explanation which links the two units together. On the whole, this suggests that agents have fragmented knowledge which only allows them to interpret and understand problem-setting poorly. This case can be properly understood as a problem of agents' bounded rationality or of high specialization.

In an alternative interpretation, the lack of explanatory links may show that agents are diverse in how they select pieces of information from the broader set of environmental stimuli, or how they are exposed to this information or are able to capture it. In this case, if an (evidence/exploratory) unit is not available to the agent, all explanatory links departing from this unit are also unavailable. This case might be considered as a problem of information scarcity. In the following, we do not distinguish exactly between these two different classes of cognitive restrictions: we focus explicitly on restrictions of the first kind (decrease in the number of available links), although in some cases they would also imply the introduction of information scarcity (e.g.: if we remove a link which is the last one connecting a unit to the schemata, that unit is not considered in the agent's problem-solving process).

As a second dimension of diversity, we model agents who are different because they might vary in the explanatory links they own. We refer to this dimension as knowledge variety. For instance, two agents may be identical in terms of knowledge amplitude,

but still diverse in the sense that the explanatory links belonging to each of them may only partially overlap. In our simulations, knowledge variety is defined in relation to a given level of knowledge amplitude in order to simplify the cases to investigate. In this way, we focus on the contribution that knowledge variety offers to a group of peers (agents displaying the same level of knowledge amplitude).

Formally, in order to define a setting in which diversity can vary in a controlled way, we model these two dimensions of diversity by introducing restrictions in the set of available knowledge constituents (explanatory links) from a complete scenario defined as the reference point. The specifications included in the next section clarify our modeling choices.

5 The Simulation Model

For the purpose of providing a common basis for the analysis of agents' performance in various communication and diversity settings, we model a simplified problem which allows two alternative representations (theories A and B), each giving access to different actions or highlighting a specific solution strategy. We define this baseline problem in terms of the sets of environmental evidence (e_i), explanatory units (a_i and b_i), and causal connections (w_{ij}), shown in Table 1 and Fig. 2. The number of units is $n = 12$, with $k = l = m = 4$ and the number of explanatory links is $g = 15$. Note that, in Fig. 2, solid lines represent positive connections between evidence and explanatory units and dashed lines represent competitive links between alternative theories. Note also that special evidence units (and their links to evidence units), cooperative links between explainers and actual weight values are omitted, for the sake of clarity.

An agent having this complete representation of the problem—hereafter called a “fully endowed” agent—will always select theory A over theory B as the result of the relaxation process.

Given this case as the reference point, we introduce agents' diversity in the model by introducing diverse partial representations of the problem. Each agent is represented by one of these partial representations, which display a differently bounded perspective for understanding the problem. Formally, each agent is represented by a subset of g explanatory links, which jointly characterize the full problem-setting.

In our simulations, subsets—and thus diversity—differ along two dimensions which we call knowledge amplitude and knowledge variety. First, subsets vary according to the number of explanatory links (shown by the h parameter). We treat the h parameter

Table 1 List of explainers in the full endowment setting

Unit of evidence	Explainers in theory A	Explainers in theory B
e_1	a_1, a_2, a_3	b_1
e_2	a_3	b_2, b_3
e_3	a_2, a_3, a_4	b_3, b_4
e_4	a_3, a_4	b_4
Tot. nr.	9	6




Fig. 2 Representation of the schemata in the full endowment setting

as a proxy of agents' knowledge amplitude, so we refer to “fully endowed agents” (for $h = g$) or “partially endowed agents” (for $h < g$).

Second, subsets of explanatory links for different partially endowed agents vary in terms of knowledge variety, by distinguishing between agents' shared versus personal explanatory links. Here, we restrict our analysis to the case of 2-person teams (agents having the same knowledge amplitude—the same value for h). We introduce knowledge variety through parameter v (with $0 \leq v \leq h$), which, for a couple of h -partial agents, measures the number of explanatory links, belonging to one agent, which differ from the explanatory links belonging to the other agent. Note that $v = 0$ if the two h -partial agents are identical (they have exactly the same explanatory links), and $v = h$ if all explanatory links of one agent are different from all the explanatory links of the other agent (to improve readability, the figures in the Sect. 6 show knowledge variety through the complement to h of v , that is, the number of explanatory links which are shared by the couple of agents—the lower this value, the higher knowledge variety).

Formally, let us define the fully endowed agent as the agent having a complete representation of the problem ($g = 15$) in terms of environmental evidence (e_i), explanatory units (a_i and b_i), and causal connections (w_{ij}). We also refer to this case as the fully endowed treatment.

The fully endowed treatment functions as a benchmark for comparing the performance of partially endowed agents, acting individually or in teams. Partially endowed agents are those who have a partial representation of the schemata shown in Fig. 2 and are modeled by means of a table of explainer which is a subset of the corresponding table in the baseline treatment. As the number of causal links available to the fully

endowed agent is $g = 15$, and a h -partial agent (for $h = 1, \dots, g - 1$) is a bounded agent having only h causal links (over the g links of the full endowment treatment), there are C_g^h possible different h -partial agents.

We measure performance in terms of the frequency with which partially endowed agents' choices coincide with those of the fully endowed agents. As the number of partially endowed agents quickly becomes very large, due to the presence of the binomial coefficient, we can offer statistics for the whole population only in extreme cases—that is, for h close to g or to zero—while in all other cases we derive our results via Monte Carlo simulations with randomly generated h -partial agents (see next section).

A final remark regarding model parameterization: we run all the instances of the model according to a choice of parameters ($d = 0.05$, $\alpha = 0.04$, $\beta = -0.06$, $s_i = 0.01$ for $j = 1, \dots, n$, $f_i = 0.1$ for $j = k + l + 1, \dots, n$) which was commonly employed in the previous literature on ECHO, and which has shown a remarkable capability for fitting data from various empirical domains. While some sensitivity analysis on these parameters showed that qualitative results do not change over a considerable parametric space, results from robustness analysis are not reported in this paper. Also, in the following section, in order to pursue the computational tractability of the problem, we restrict ourselves to the case of the smallest possible team ($p = 2$). Lastly, we treat communication as an on-off variable ($\delta = 0$ or $\delta = 0.5$) for all simulations except the last one, where we focus on the effects of different communication strengths, increasing δ to 1 and 1.5.

6 Results

We organize our results into five subsections. In the first, we show the performance of partially endowed agents as individual problem-solvers (no communication occurs), and explore the influence on the performance of various levels of endowment. These results are used as a benchmark in the subsequent subsections, in which we study how performance is affected by the introduction of team communication.

In Sect 6.2, we introduce 2-person teams made up of one fully endowed agent and one partially endowed agent as a way to study the influence of increasing differences in knowledge amplitude within a team.

In Sect. 6.3, we study 2-person teams made up of two partially endowed agents having the same degree of knowledge amplitude. Here, the results are compared with the individual benchmark (Sect. 6.1) in order to assess to what extent communication can make up for partiality. In all these three treatments, we measure performance by varying agents' endowment within the whole domain of the h parameter. Treatment three is also examined in the two remaining subsections.

Section 6.4 examines the interplay between different sources of diversity in 2-person teams: we study how performance is jointly affected by bounds in agents' explanatory ability (that is, $h < g$) and the variety of agents' knowledge (different elements in the h causal links). Finally, the last subsection extends the previous analysis, providing a simple sensitivity analysis of the communication strength parameter.

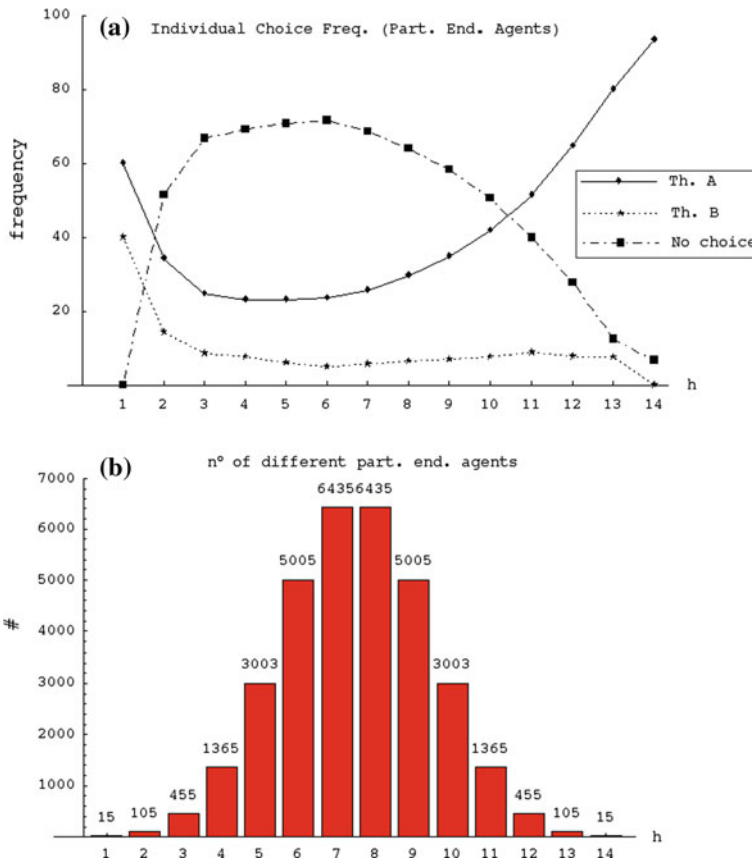


Fig. 3 **a** Frequencies of the correct interpretation (“theory A”), wrong one (“theory B”) and of reaching an inconclusive outcome (inability of selecting between the two theories), computed for h -partial agents (according to the various values of h) and **b** frequencies of the possible h -partial agents, computed according to the binomial coefficient C_{15}^h , for $h = 1, \dots, 14$

6.1 Agents as Individual Problem-Solvers: Partial Representations of the Problem and Performance

Figure 3a shows the outcome on individual performance resulting from introducing limits in knowledge amplitude. Note that, since agents do not communicate in this case, we are not investigating issues of diversity among interacting agents: we are interested here in varying knowledge amplitude only as a means to understand to what extent, on an individual basis, performance is affected by the diversity modeled as deterioration of knowledge regarding a problem representation. Data are organized in order to distinguish poorly endowed agents (low values of the h parameter) from “almost fully” endowed agents (high values); frequencies are computed over the whole populations of possible h -partial agents, whose sizes are shown in Fig. 3b.

Recalling that a fully endowed agent always solves the problem in terms of theory A, a decline in the capability of correctly interpreting the problem is clearly visible if one moves from the right to the left side of the plot: agents who are more and more bounded in the representation of the problem select theory A at a decreasing rate. Interestingly, selection of theory B does not increase correspondingly. On the contrary, the decrease in performance is almost entirely due to the emergence of confusion as the result of agents' efforts to interpret the problem: they select outcomes which are inconclusive, as they do not highlight any winner between the two available theories. Formally, the system converges to fixed points, in which the values of units s_j do not allow us clearly to distinguish a winning theory (the theory whose units are all positive) from a losing one (the theory whose units are all negative).

Note that, in the case of very poor problem settings (for $h \leq 5$), agents tend to get less stuck into inconclusive outcomes, as it becomes more and more probable that the few remaining links provide unambiguous, strong support for one theory over the other one.

6.2 Performance in Diverse Teams: The Role of Knowledge Amplitude

We first address the topic of diversity by studying the influence of increasing differences in knowledge amplitude in teams. For the sake of simplicity, we restrict our analysis to the case of 2-person teams composed of one fully endowed agent and one partially endowed agent, and model increases in diversity by decreasing the number of links available to the latter agent. A widespread practice in organizations is to group workers having vast knowledge with novices. Through this simulation, we are interested in assessing whether poorly endowed agents might benefit by interacting with fully endowed colleagues. Accordingly, results should indicate whether the latter agents' performance may be negatively affected.

Figure 4 shows the observed frequencies of performance, respectively for (a) the partially endowed agent and (b) the fully endowed agent. Figure 4b clearly shows that the effect of communicating with a bounded agent has no effect at all on the fully endowed agent's performance. Instead, communication helps the partially endowed agent to avoid the trap represented by the wrong theory (B), which is never selected. Nevertheless, the partially endowed agent's performance is still far from the levels of the fully endowed one. In fact, as in the baseline treatment, when his knowledge was very poor, the frequency of inconclusive outcomes is still higher than the frequency of the correct interpretation (theory A). Despite this, an overall look at the results shows that communication has a positive effect on teams composed of an experienced worker and a novice at any level of h , as shown in Fig. 5.

6.3 Teams with Two Partially Endowed Agents

Figure 6 introduces team communication between couples of partially endowed agents. We restrict analysis to the symmetric case and assume that a h -partial team is composed of two h -partial agents (note here that we do not control for knowledge variety, which is examined in the next subsection). Since the sizes of the populations of the




Fig. 4 Frequencies of the correct interpretation ("theory A"), of the wrong one ("theory B") and of reaching an inconclusive outcome, respectively, for **a** a partially endowed agent and **b** a fully endowed agent, communicating in a 2-person team

possible h -partial teams, as shown in Fig. 6b, are large, frequencies depicted in Fig. 6a were collected by random sampling of 20,000 h -partial agents randomly grouped in 10,000 teams (data for the two agents are pooled in the analysis). From Figs. 6a and 7, which compares the observed frequency of the correct interpretation over the three treatments (partial agent alone, partial agent matched with a full agent, partial agent matched with another similarly partial agent), we can derive the following considerations: (i) communicating in a team with a fully endowed agent improves the performance of the partially endowed agent at all levels of h . In this respect, the partially endowed agent will always prefer to team with the fully endowed agent than with similarly partially endowed ones; (ii) when two partially endowed teammates

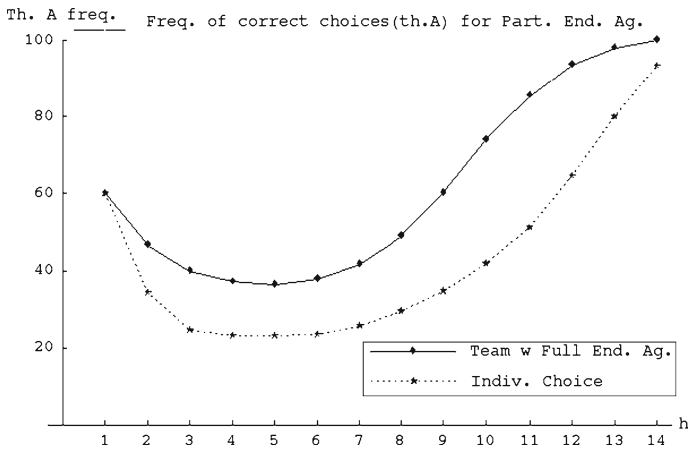


Fig. 5 Frequencies of the correct interpretation (“theory A”) for a partially endowed agent with and without team communication with a fully endowed partner

form a group, performance follows a less straightforward trend. When peers’ problem representation, albeit limited, still gives a rough picture of the situation (intermediate levels of h), team interaction can improve agents’ performance. Instead, as their limits in representing the problem increase, they may reach a point after which their performance as individuals seems more effective, since the frequency of inconclusive outcomes in the communication treatment becomes larger.

An intuitive interpretation for these results is that, when agents can only represent the problem they are facing poorly, communication may drive them into a decision trap, making them incapable of reconciling suggestions coming from their teammate with their own similarly fragmented representation. One explanation for the poor performance of communication may lie in the observation that, the more limited agents’ problem representations are, the higher the probability that their schemata will have very few elements in common, thus most probably resulting in diverse patterns of activation of the units for both agents. Consequently, in this case we expect that communication will increase confusion, since agents exchange beliefs over very diverse sets of activation of units. Consider, for example, Fig. 8a and b, in which two agents displaying a rich representation of the problem are able to improve their performance through communication. Their schemata is made up of 12 of the 15 links of the complete representation. Although the 3 links they miss are different, there are still 9 links which are shared. Conversely, Fig. 8c and d show two agents having a very poor representation of the schemata, with only 3 links. Through communication, the performance of agent 1 does not improve (with reference to the fully endowed agent, who chooses theory A), while agent 2, who was correct first, changes his mind and becomes confused, worsening his performance. In order to examine better the relation between agents’ completeness of problem representations and their overlap, we integrate another dimension of diversity into the analysis: knowledge variety.

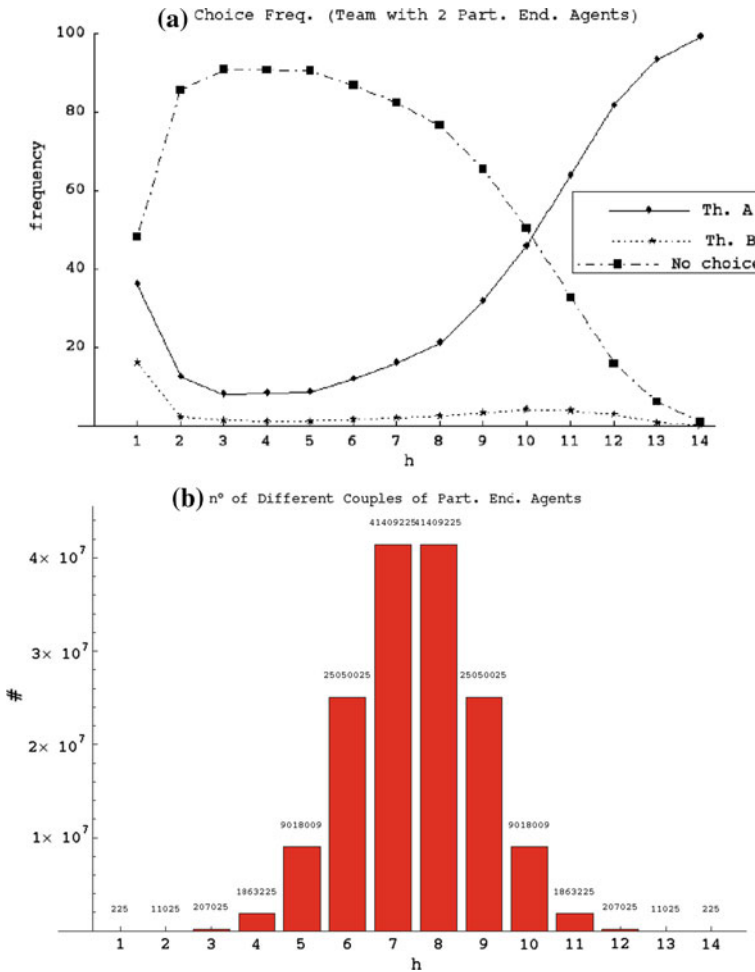


Fig. 6 **a** Frequencies of the correct interpretation (“theory A”), wrong interpretation (“theory B”) and of reaching an inconclusive outcome (inability of discriminating between the two alternative interpretations), computed for a h -partial agent communicating in a 2-person team with another h -partial agent, **b** frequencies of the possible couples of h -partial agents, computed according to the formula $(C_{15}^h)^2$, for $h = 1, \dots, 14$. Frequencies in **a** are computed with respect to random samples of 10.000 couples of h -partial agents

6.4 The Interplay Between Diversity Dimensions: Knowledge Amplitude versus Knowledge Variety in 2-Agent Teams

We measure knowledge variety for a given couple of h -partial agents as the complement to h of the number of common causal links shared by the couple. In Table 2, which counts the possible h -partial teams with respect to h and to the number of shared causal links, the level of variety increases moving from the right to the left side of the table.

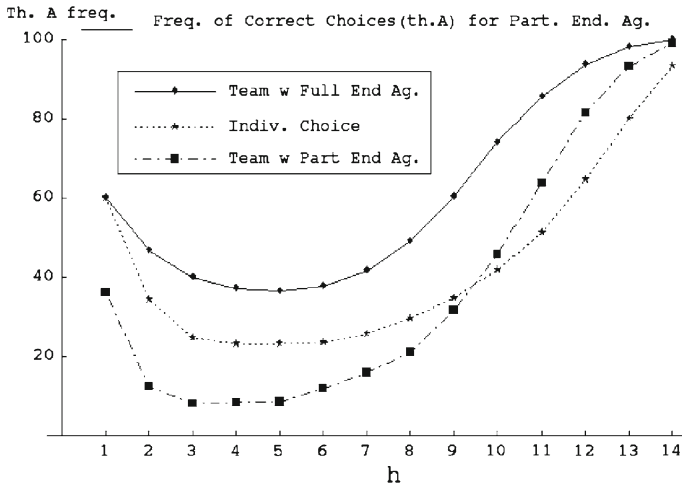


Fig. 7 Frequencies of the correct interpretation (“theory A”) for a partially endowed agent under three different treatments: individual behavior, communication with a fully endowed partner, communication with a partially endowed partner

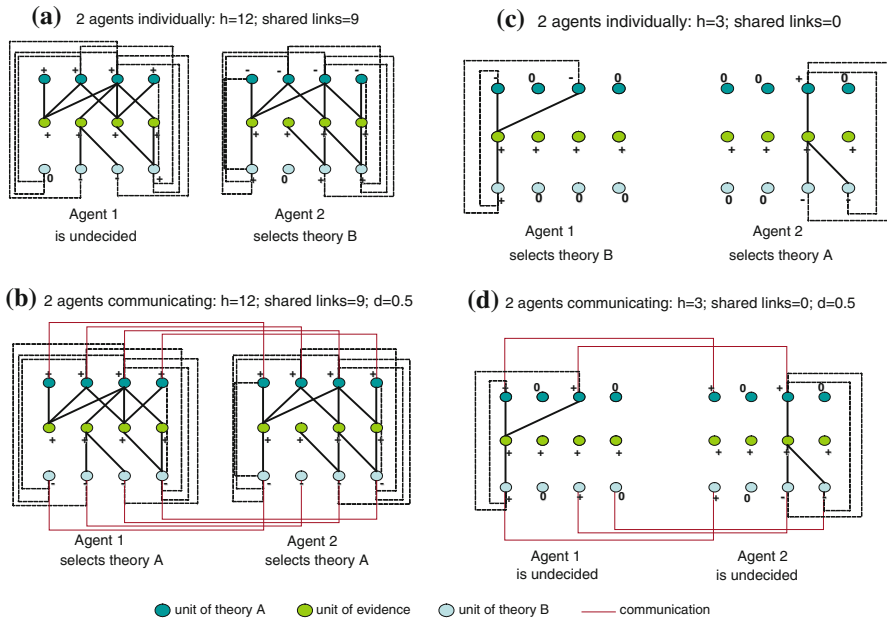


Fig. 8 Representation of the schemata of partially endowed agents and of their final activation individually and communicating

Figure 9 shows the frequencies of selecting the correct interpretation (theory A) in an h -partial team for various levels of knowledge amplitude (h) and variety (measured in the horizontal axis). Frequencies are computed over the whole corresponding team population if, according to Table 2, its size is less than 10,000; otherwise frequencies

refer to 10,000 random couples. The points in the plot refer to the case in which the number of shared links is equal to h , thus representing teams of 2 identical agents. Note that, in this case, outcomes correspond exactly to the performance of individual agents previously shown in Fig. 3. In the proximity of each point, two lines depart toward the left side of the plot: the solid lines represent the frequency of the correct interpretation under team communication for increasing levels of variety in the couple of agents (data are pooled for both of them); the dashed lines represent how these agents would perform without communication (these frequencies are slightly different from the values at each corresponding right-end point, because they were computed over different random samples).

Overall, the results shown in Fig. 9 shed light on the interplay between three intertwined elements: team communication, agents' knowledge amplitude in problem representation, and knowledge variety. For high levels of knowledge endowment (right side of Fig. 9), communication always improves individual outcomes and, the more various knowledge agents have, the higher their probability of inferring the correct interpretation. This trend shows that, when agents have an almost complete representation of the problem, they can benefit most from interaction, because they can enrich their view of the problem with aspects which they had ignored or were not able to explain. In short, communication helps them to solve the puzzle. The circulation of beliefs on problem representation is an effective mechanism for completing agents' perspectives, as long as the individual representations have quite a high number of common elements. In fact, we do observe that, for intermediate levels of knowledge endowment (center of Fig. 9), agents' performance still benefits from communication, although in this case there is an optimal amount of variety guaranteeing the highest performance levels. After these, performance declines to a point at which extreme levels of variety result in worse performance than in the case of independent individual problem-solving. Lastly, for lower levels of knowledge endowment (left side of Fig. 9), team communication always results in worse outcomes than the individual case, and increasing variety further hampers performance.

Note that, for $8 \leq h \leq 13$ and no variety (the number of shared links is exactly h), the solid lines do not start from the points representing the performance of 2 identical agents not communicating with each other. This means that communication improves performance also in the case of identical agents. This happens because, in the no communication treatment, agents sometimes fail to recognize the correct theory because of only one unit which has discordant activation (all the units of theory A are positive and those of theory B are negative, with the exception of one which is slightly positive). Thus, in the communication treatment, the activation values of A increase and those of B decrease; in addition, the relations among units make this effect stronger. As a result, the only unit which was previously discordant turns out being null or negative, which makes it possible to select the correct theory, thus improving performance.

This evidence shows that groups can benefit most from diverse members—in terms of variability of knowledge—when they have a very good representation of the problem. Conversely, when their understanding of it is poor, diversity results in declining performance.

Table 2 Frequencies of the possible couples of h -partial agents, computed distinguishing for the level of endowment (h) and for the variety among the h features available to each of the two agents (that can be measured in terms of the complement to h of the number of causal links that they share)

h	Shared links														
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	210	15	-	-	-	-	-	-	-	-	-	-	-	-	-
2	8,190	2,730	105	-	-	-	-	-	-	-	-	-	-	-	-
3	100,100	90,090	16,380	455	-	-	-	-	-	-	-	-	-	-	-
4	450,450	900,900	450,450	60,060	1,365	-	-	-	-	-	-	-	-	-	-
5	756,756	3,153,150	3,603,600	1,351,350	150,150	3,003	-	-	-	-	-	-	-	-	-
6	420,420	3,783,780	9,459,450	8,408,400	2,702,700	270,270	5,005	-	-	-	-	-	-	-	-
7	51,480	1,261,260	7,567,560	15,765,750	12,612,600	3,783,780	360,360	6,435	-	-	-	-	-	-	-
8	-	51,480	1,261,260	7,567,560	15,765,750	12,612,600	3,783,780	360,360	6,435	-	-	-	-	-	-
9	-	-	-	420,420	3,783,780	9,459,450	8,408,400	2,702,700	270,270	5,005	-	-	-	-	-
10	-	-	-	-	-	756,756	3,153,150	3,603,600	1,351,350	150,150	3,003	-	-	-	-
11	-	-	-	-	-	-	-	450,450	900,900	450,450	60,060	1,365	-	-	-
12	-	-	-	-	-	-	-	-	-	100,100	90,090	16,380	455	-	-
13	-	-	-	-	-	-	-	-	-	-	-	8,190	2,730	105	-
14	-	-	-	-	-	-	-	-	-	-	-	-	-	210	15

Note that the elements a_{ij} (for $i = 1, \dots, 14$) correspond to the frequencies of the populations of h -partial agents shown in Fig. 3a, and horizontal sums corresponds to the frequencies of the population of h -partial teams shown in Fig. 6b. Also note that from the lower to the upper bound of the Table we move towards agents more bounded in their knowledge, and that from right to left we move, ceteris paribus, toward agents with less shared links, thus displaying more variety in their knowledge

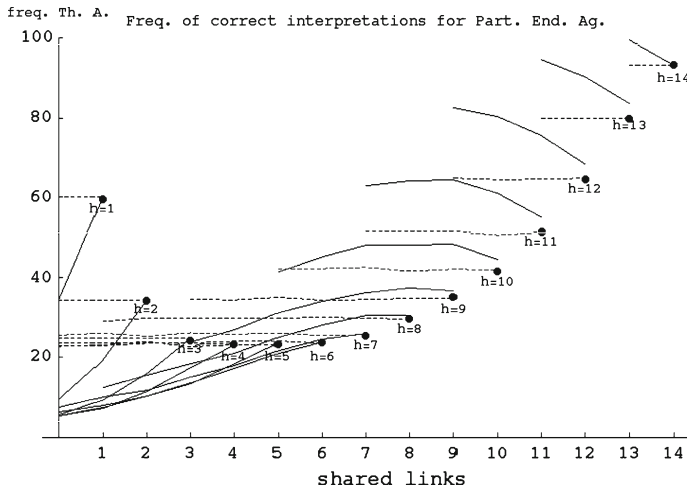


Fig. 9 Frequencies of the correct interpretation (theory A) measured for h -partial teams under various level of endowment (h) and variety (measured on the horizontal axis, fewer shared links corresponds to a higher variety in the team)

6.5 The Role of Communication Strength

Results from the previous subsections were obtained with a degree of communication strength which presupposes moderate interaction among agents. In this subsection, we perform a discrete sensitivity test by doubling and tripling the value of the communication strength parameter. Recalling how the schemata is initialized, these three values correspond, respectively, to the case in which communication has a lower, similar or higher influence on the weights of the communication matrices, in comparison with the default intensity of the causal relationships of the individual weight matrices. We restrict analysis to these values, arguing that this choice is consistent with the aim of preserving enough intelligibility in the graphical presentation of results without losing depth in analysis, since outcomes for lower or higher communication strengths can easily be derived from the analysis carried out at the levels we focus on in the following. Results are shown in Fig. 10, in which the solid lines are arranged so that the thicker lines correspond to higher communication strengths. Due to the multiple series of data, in order to improve readability, we show the results only for selected h values.

We arrange our results distinguishing between low/high levels of endowment. For low levels (left side of Fig. 10, for $h \leq 8$), sensitivity analysis clearly confirms the negative contribution of communication: all communication levels produce a worse performance than individual outcomes. Also, communicating more does not help agents who have a poor understanding of the problem setting: there is a clear-cut and negative relationship linking performance and communication strength, resulting in communication mishaps which are particularly evident for the highest δ value. This means that teams of agents having very poor representations and over-discussing the details of




Fig. 10 Frequencies of the correct interpretation (theory A) measured for h -partial teams under various level of endowment (h) and variety (measured on the horizontal axis, fewer shared links corresponds to a higher variety in the team), for different degrees of communication strengths: the *thinnest (solid)* line corresponds to the initial treatment displayed in Fig. 9 (the lowest, with $\delta = 0.5$), the *thickest* one to a highest communication level ($\delta = 1.5$), while the intermediate *thickness* corresponds to an intermediate strength level ($\delta = 1$)

the problem are more likely to distort their relevance in the general picture, and to get stuck at a point in which they cannot see which choice should be made.

For higher levels of endowment, ($h \geq 10$) findings are less straightforward and may be summarized as follows. First, there is a clear non-monotonic relationship between communication strengths and performance: although we did not perform a complete exploration of the strength space in order to find an optimal value, it is clear that performance corresponding to the intermediate communication value is always better than either the lower or the higher strength levels. Second, as endowment increases, sensitivity to changes in communication strength decreases, since differences in performance improvements become smaller. Third, the interplay between communication and variety in agents' knowledge seems to be particularly unanticipated. In fact, while performance differentials for various communication strengths are negligible in the case of very diverse agents, they are very large as agents become more and more similar. The non-linearity of the phenomenon is startling: while doubling the communication strength results in relatively limited improvements in performance, when moving from $\delta = 1$ to $\delta = 1.5$, we observe a severe breakdown in performance. Lastly, higher communication strengths may affect behavior even in teams of identical agents: while for $\delta = 0.5$ communication did not result in any observable difference with respect to individual behavior, as the strength of communication increased, identical agents changed their behavior when interacting in a team. More precisely, this results in improvements at the intermediate communication level; at the highest level, it corresponds to a considerable deterioration in agents' problem-solving abilities.

7 Discussion and Conclusions

In this paper we examined the influence of communication in teams of diverse problem-solvers. We modeled diversity as differences in problem representation. Extant literature addresses the relation between team diversity and performance through models that do not allow us to control for various determinants and sources of diversity in teams of heterogeneous agents. Instead, our model defines diversity along two different dimensions: knowledge amplitude and knowledge variety, thus opening the door to the exploration of the effects of distinct features and levels of diversity on the performance of problem-solving teams.

We studied interactions among diverse peers, who were agents having the same knowledge amplitude, but who may be diverse in terms of knowledge variety. Likewise, we provided results regarding agents' interactions in teams composed of poorly endowed and fully endowed problem-solvers as a way of examining the role of diverse levels of knowledge amplitude in problem-solving performance.

Our results allow us to derive some implications on group composition and interaction in firms and organizations. Our main findings show that teams are not always effective in tackling problem-solving in organizations.

This is certainly not the case of teams pairing more endowed workers with partners having a poor understanding of the problem. The former are not diverted by their teammate's partial understanding, since they can lead them to solutions which are better, without recording significant losses in their performance. This evidence supports earlier findings on the benefits of the expert-novice relationship ([Azmitia 1988](#); [Daiute and Dalton 1993](#); [Rao and Joon 1995](#); [John-Steiner 2000](#); [Barron 2003](#)). It is also interesting for studies on the roles of people in communication. Our model can control for competence and information exposure on the problem and, interestingly enough, it shows that in a team of peers, when competence on a problem by one of the team members is acknowledged and accepted, it is beneficial for performance.

Conversely, if we address the topic of communication effectiveness in teams of peers, results may vary. In particular, communication may have a negative influence when the understanding of the situation is very fragmented. We showed that, when peers display a very limited understanding of the problem, their propensity for failure increases. When agents do not share a considerable overlap in their problem representation, their perspectives on the world diverge and communication does not find a common basis for beneficial interaction. In this scenario, communication may eventually be more troublesome than helpful. Moreover, adding variety into these poor interpretations makes the situation even worse: agents increasingly confuse each other's interpretation as they become more diverse. These results support the "pessimistic view" on diversity ([Mannix and Neale 2005](#)), showing the negative influence of diversity on performance. Nevertheless, our evidence does provide an original contribution to this line of research. Our model shows that, when an agent presides over part of the issue but is unable to provide a sufficient explanation to the whole problem, he appears to be unpersuasive and is not considered. In this sense, highly diverse people do not share any knowledge, nor do they acknowledge any competence to others whose knowledge is too poor for the problem at stake. In such a situation, discussing the terms of the problem is more confusing than beneficial, as none of those present

is able to make their view on the problem sound convincing. This shows that it is not only important to clarify attributes defining diversity and to be able to measure the amount of diversity and its distribution among team members, but also to be able to assess members' potential contribution to problem solution. When the latter is too low, then collaboration among teams may end up inconclusive rather than beneficial. Thus, low performance may not only be due to diversity or communication difficulties, but also to poor contributions to problem solution.

Communication helps when peers have a vast knowledge domain which allows them to identify the majority of the relevant information and to define consistent and coherent explanations of the evidence. Interaction allows them to complement each other's knowledge in an effective way, and increases in knowledge variety can also further improve team performance.

Notwithstanding these results, we also show that interacting too much may change the overall influence of communication. This is not the case of teams revealing high variety in knowledge, meaning that talking more or more intensely with people who are well-read in domains that are different from our own is enriching and eventually improves performance. Diversity works when people are highly competent in different domains. Intense communication then allows them to learn reciprocally.

However, very similar agents are extremely sensitive to higher-than-optimal communication strengths, as their performance declines considerably. This shows that increasing communication intensity does not always help in sorting out the interpretation of a problem setting. As a matter of fact, quite the reverse often happens, since agents tend to get trapped in undecidable outcomes as the result of over-communication. Interesting enough, the opposite case of interacting too little is less disruptive, as lower-than-optimal communication strengths seem to have a less significant influence on performance. As an implication, these results indicate that there is a negative relationship between the intensity of communication (measured either as the absolute amount of interactions or in terms of frequencies) among similar team members and performance. This instance represents the typical homogeneous group sketched by the social comparison theory (Festinger 1954), in which people are similar in order to respond to the need to compare themselves to their own kind constantly. In such teams, agents feel pressure toward uniformity as they try to remove any discrepancy in opinions. As a result, low communication levels are enough to conform opinions, whereas high communication drives the need for uniformity too far, so that people are confused about what they should conform to.

Our results are limited in various ways. First of all, they rely on a specific, albeit commonly reported choice of model parameters and examine the specific instance of a problem setting which has an abstract and acceptable although arbitrary structure. The reasons for starting from this unique and ad hoc instance may be justified when we take into account the fact that: (i) we needed to start from a full representation of problem setting which relaxed towards one theory, something that is not granted by generating lists of explainers at random (like the one shown in Fig. 2), since the corresponding fixed point might be inconclusive; (ii) computational parsimony suggested restricting analysis to one structure (as opposed to averaging results for a series of different instances) and to limit the number of units involved ($n = 12$). Similar reasons suggested limiting our analysis to the case of 2-agent teams and two competing

theories. Lastly, we explored knowledge variety only in the case of similarly endowed agents (peers); more complex cases where both amplitude and variety are different across team members still need to be addressed.

The current model does not pretend to be complete, and several extensions could be made. We offer it with a double aim: first, to examine the contribution of diversity to team problem-solving, thus providing deeper understanding of the dynamics of collective problem-solving; second, to explore models of decision-making which are detached from the pattern-matching basis or from bounded evaluations of consequences.

Acknowledgments We gratefully acknowledge valuable suggestions and comments received on earlier versions of this paper from Paolo Collini, Luigi Marengo, James March, Ricardo Marques Pereira, Alessandro Narduzzo, Paul Thagard, Massimo Warglien, Enrico Zaninotto and one anonymous referee. We have benefited from constructive criticisms by ROCK Group members during work in progress, and by participants attending to our presentations held at Scancor, Stanford University, on May 3, 2007, at the Computational Methods for Modelling and Learning in the Social and Human Sciences Bilingual Conference at ENST, Brest, France, on May 11, 2007, and at the ESHIA/WEHIA International Conference on Economic Science with Heterogeneous Interacting Agents, Fairfax, Virginia, on June 17–20, 2007. The usual disclaimer applies. We received partial funding for this research from the Italian Ministero dell'Istruzione, dell'Università e della Ricerca under the projects FIRB03 and PRIN05. One of us was supported by the Professor Claudio Demattè Memorial Grant during part of this research.

References

- Ahn TK, Isaac RM, Salmon TC (2008) Endogenous group formation. *J Public Econ Theory* 10:171–194. doi:[10.1111/j.1467-9779.2008.00357.x](https://doi.org/10.1111/j.1467-9779.2008.00357.x)
- Ancona D, Caldwell D (1992) Demography and design: predictors of new product team performance. *Organ Sci* 3:321–341. doi:[10.1287/orsc.3.3.321](https://doi.org/10.1287/orsc.3.3.321)
- Axelrod R (1997) The complexity of cooperation. Agent-based models for competition and collaboration. Princeton University Press, Princeton
- Azmitia M (1988) Peer interaction and problem solving: when are two heads better than one? *Child Dev* 59:87–96. doi:[10.1080/17405620701860165](https://doi.org/10.1080/17405620701860165)
- Bantel K, Jackson S (1989) Top management and innovations in banking: does the composition of the team make a difference. *Strateg Manage J* 10:107–124. doi:[10.1002/smj.4250100709](https://doi.org/10.1002/smj.4250100709)
- Barron B (2003) When smart groups fail. *J Learn Sci* 12:307–359. doi:[10.1207/S15327809JLS1203_1](https://doi.org/10.1207/S15327809JLS1203_1)
- Carpenter M (2002) The implications of strategy and social context for the relationship between top management team heterogeneity and firm performance. *Strateg Manage J* 23:275–284. doi:[10.1002/smj.226](https://doi.org/10.1002/smj.226)
- Chatman JA, Flynn FJ (2001) The Influence of Demographic Heterogeneity on the Emergence and Consequences of Cooperative Norms in Work Teams. *Acad Manag J* 44:956–974.
- Daiute C, Dalton B (1993) Collaboration between children learning to write: can novices be masters? *Cognit Instr* 10:281–333. doi:[10.1207/s1532690xci1004_1](https://doi.org/10.1207/s1532690xci1004_1)
- Deetz S (2001) Conceptual foundations. In: Putnam L, Jablin F (eds) *The new handbook of organizational communication*. Sage, Thousand Oaks, pp 3–46
- Eliasmith C, Thagard P (1997) Waves, particles, and explanatory coherence. *Br J Philos Sci* 48:1–19. doi:[10.1093/bjps/48.1.1](https://doi.org/10.1093/bjps/48.1.1)
- Festinger L (1954) A theory of social comparison processes. *Hum Relat* 7:117–140. doi:[10.1177/001872675400700202](https://doi.org/10.1177/001872675400700202)
- Frigotto ML, Rossi A (2006) Modeling strategic decision making in novel contexts: a constraint satisfaction approach. Quaderno DISA n. 117, Department of Computer and Management Sciences, Trento University, Italy.
- Geard NL, Bullock S (2008) Group formation and social evolution: a computational model. In: Bullock S, Noble J, Watson R, Bedau M (eds) *Artificial life XI: proceedings of the eleventh international conference on the simulation and synthesis of living systems*. MIT Press, Cambridge, MA, pp 197–203

- Guzzo R, Dickson M (1996) Teams in organizations: recent research on performance and effectiveness. *Annu Rev Psychol* 47:307–338. doi:[10.1146/annurev.psych.47.1.307](https://doi.org/10.1146/annurev.psych.47.1.307)
- Hoffman L (1959) Homogeneity and member personality and its effect on group problem solving. *J Abnorm Soc Psych* 58:27–32. doi:[10.1037/h0043499](https://doi.org/10.1037/h0043499)
- Hoffman L, Maier N (1961) Quality and acceptance of problem solutions by members of homogeneous and heterogeneous groups. *J Abnorm Soc Psych* 62:401–407. doi:[10.1037/h0044025](https://doi.org/10.1037/h0044025)
- Hong L, Page S (2001) Problem solving by heterogeneous agents. *J Econ Theory* 97:123–163. doi:[10.1006/jeth.2000.2709](https://doi.org/10.1006/jeth.2000.2709)
- Hong L, Page S (2004) Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *PNAS* 101:16385–16389. doi:[10.1073/pnas.0403723101](https://doi.org/10.1073/pnas.0403723101)
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *PNAS* 79:2554–2558. doi:[10.1073/pnas.79.8.2554](https://doi.org/10.1073/pnas.79.8.2554)
- Huck S, Rey-Biel P (2006) Endogenous leadership in teams. *J Inst Theor Econ* 162:253–261. doi:[10.1628/093245606777583495](https://doi.org/10.1628/093245606777583495)
- Hutchins E (1995) *Cognition in the wild*. The MIT Press, Cambridge
- Ilgel DR (1999) Teams embedded in organizations: some implications. *Am Psychol* 54:129–139. doi:[10.1037/0003-066X.54.2.129](https://doi.org/10.1037/0003-066X.54.2.129)
- Jackson SE, Joshi A, Erhardt NL (2003) Recent research on team and organizational diversity: SWOT analysis and implications. *J Manage* 29:801–830. doi:[10.1016/S0149-2063_03_00080-1](https://doi.org/10.1016/S0149-2063_03_00080-1)
- Jackson SE, Stone VK, Alvarez EB (1992) Socialization amidst diversity: the impact of demographics on work team oldtimers and newcomers. In: Cummings LL, Staw W (eds) *Research in organizational behavior*, vol 15. JAI Press, Greenwich, pp 45–109
- Jehn KA, Northcraft GB, Neale MA (1999) Why differences make a difference: A field study of diversity, conflict, and performance in workgroups. *Adm Sci Quart* 44:741–763
- John-Steiner V (2000) *Creative collaboration*. Oxford University Press, New York
- Keil FC (2006) Explanation and understanding. *Annu Rev Psychol* 57:227–254. doi:[10.1146/annurev.psych.57.102904.190100](https://doi.org/10.1146/annurev.psych.57.102904.190100)
- Keser C, Montmarquette C (2007) Voluntary teaming and effort. Discussion Papers of DIW Berlin 745, DIW Berlin, German Institute for Economic Research
- Kilduff M, Angelmar R, Mehra A (2000) Top-management team diversity and firm performance: examining the role of cognitions. *Organ Sci* 11:21–34. doi:[10.1287/orsc.11.1.21.12569](https://doi.org/10.1287/orsc.11.1.21.12569)
- Kocher M, Strauß S, Sutter M (2006) Individual or team decision-making—causes and consequences of self-selection. *Game Econ Behav* 56:259–270. doi:[10.1016/j.geb.2005.08.009](https://doi.org/10.1016/j.geb.2005.08.009)
- Kozlowski SWJ, Klein KJ (2000) A multilevel approach to theory and research in organizations: contextual, temporal, and emergent processes. In: Klein KJ, Kozlowski SWJ (eds) *Multilevel theory, research, and methods in organizations: foundations, extensions, and new directions*. San Francisco, Jossey-Bass, pp 3–90
- Mannix E, Neale MA (2005) What differences make a difference?. *Psychol Public Interest* 6:31–55. doi:[10.1111/j.1529-1006.2005.00022.x](https://doi.org/10.1111/j.1529-1006.2005.00022.x)
- March JG (1975) *The pursuit of organizational intelligence*. Blackwell, Malden
- March JG (1994) *A primer on decision making: how decisions happen*. Free Press, New York
- March JG, Sproull LS, Tamuz M (1991) Learning from samples of one or fewer. *Organ Sci* 2:1–13. doi:[10.1287/orsc.2.1.1](https://doi.org/10.1287/orsc.2.1.1)
- Marchiori D, Warglien M (2005) Constructing shared interpretations in a team of intelligent agents: the effects of communication intensity and structure. In: Terano T (ed) *Agent-based simulation: from modeling methodologies to real-world applications*. Post-proceedings of the 3rd international workshop on agent-based approaches in economic and social complex systems 2004. Springer, Berlin, pp 58–71
- Milliken FJ, Martins LL (1996) Searching for Common Threads: Understanding the Multiple Effects of Diversity in Organizational Groups. *Acad Manage Rev* 21:402–433.
- Morgeson FP, Hofmann DA (1999) The structure and function of collective constructs: implications for multilevel research and theory development. *Acad Manage Rev* 24:249–265. doi:[10.2307/259081](https://doi.org/10.2307/259081)
- Morris M, Moore P (2000) The lessons we (don't) learn: counterfactual thinking and organizational accountability after a close call. *Adm Sci Quart* 45:737–765
- Nemeth CJ (1986) Differential contributions of majority and minority influence. *Psychol Rev* 93:23–32. doi:[10.1037/0033-295X.93.1.23](https://doi.org/10.1037/0033-295X.93.1.23)
- Newell A, Simon HA (1972) *Human problem solving*. Prentice Hall, Englewood Cliffs

- Nowak G, Thagard P (1992a) Copernicus, ptolemy, and explanatory coherence. In: Giere R (ed) *Cognitive models of science*, vol 15. University of Minnesota Press, Minneapolis, pp 274–309
- Nowak G, Thagard P (1992b) Newton, descartes, and explanatory coherence. In: Duschl R, Hamilton RJ (eds) *Philosophy of science, cognitive psychology and educational theory and practice*. SUNY Press, Albany, pp 69–115
- Page T, Putterman L, Unel B (2005) Voluntary association in public goods experiments: reciprocity, mimicry and efficiency. *Econ J* 115:1032–1053. doi:[10.1111/j.1468-0297.2005.01031.x](https://doi.org/10.1111/j.1468-0297.2005.01031.x)
- Patriotta G (2003) Sensemaking on the shop floor: narratives of knowledge in organizations. *J Manage Stud* 40(2):349–375. doi:[10.1111/1467-6486.00343](https://doi.org/10.1111/1467-6486.00343)
- Pelled LH, Eisenhardt KM, Xin KR (1999) Exploring the Black Box: An Analysis of Work Group Diversity, Conflict, and Performance. *Adm Sci Quart* 44:1–28.
- Pitcher P, Smith AD (2001) Top management team heterogeneity: personality, power, and proxies. *Organ Sci* 12:1–18. doi:[10.1287/orsc.12.1.1.10120](https://doi.org/10.1287/orsc.12.1.1.10120)
- Rao HR, Joon MA (1995) The effect of team composition on decision scheme, information search, and perceived complexity. *J Organ Comput Elect Com* 5:1–20. doi:[10.1080/10919399509540236](https://doi.org/10.1080/10919399509540236)
- Rumelhart DE, Smolensky P, McClelland JL, Hinton GE (1986) Schemata and sequential thought processes in PDP models. In: Rumelhart DE, McClelland JL, The PDP Research Group (eds) *Parallel distributed processing: explorations in the microstructure of cognition*, vol 2. MIT Press, Cambridge, pp 7–57
- Schneider B, Smith DB, Sipe WP (2000) Personnel selection psychology: multilevel considerations. In: Klein KJ, Kozlowski SWJ (eds) *Multilevel theory, research, and methods in organizations: foundations, extensions, and new directions*. San Francisco, Jossey-Bass, pp 91–120
- Simon HA (1991) Bounded rationality and organizational learning. *Organ Sci* 2:125–134. doi:[10.1287/orsc.2.1.125](https://doi.org/10.1287/orsc.2.1.125)
- Thagard P (1989) Explanatory coherence. *Behav Brain Sci* 12:435–467
- Thagard P (1992a) Adversarial problem solving: modeling an opponent using explanatory coherence. *Cogn Sci* 16:123–149. doi:[10.1207/s15516709cog1601_4](https://doi.org/10.1207/s15516709cog1601_4)
- Thagard P (1992b) *Conceptual revolutions*. Princeton University Press, Princeton
- Thagard P (1998a) Ulcers and bacteria I: discovery and acceptance. *Stud Hist Philos Biol & Biomed Sci* 29:107–136. doi:[10.1016/S1369-8486\(98\)00006-5](https://doi.org/10.1016/S1369-8486(98)00006-5)
- Thagard P (1998b) Explaining disease: correlations, causes, and mechanisms. *Mind Mach* 8:61–78. doi:[10.1023/A:1008286314688](https://doi.org/10.1023/A:1008286314688)
- Thagard P (2000) *Coherence in thought and action*. MIT Press, Cambridge
- Thagard P (2004) Causal inference in legal decision making: explanatory coherence versus bayesian networks. *Appl Artif Intell* 18:231–249. doi:[10.1080/08839510490279861](https://doi.org/10.1080/08839510490279861)
- Thagard P, Kroon FW (2006) Emotional consensus in group decision making. *Mind Soc* 5:85–104. doi:[10.1007/s11299-006-0011-5](https://doi.org/10.1007/s11299-006-0011-5)
- Trenholm S (1986) *Human communication theory*. Prentice-Hall, Englewood Cliffs
- Triandis H, Hall E, Ewen R (1965) Member heterogeneity and dyadic creativity. *Hum Relat* 18:33–55. doi:[10.1177/001872676501800104](https://doi.org/10.1177/001872676501800104)
- Watson WE, Kumar K, Michaelsen LK (1993) Cultural diversity's impact on interaction process and performance: comparing homogeneous and diverse task groups. *Acad Manag J* 36:590–602
- Webber SS, Donahue LM (2001) Impact of highly and less job-related diversity on work group cohesion and performance: a meta-analysis. *J Manag* 27:141–162. doi:[10.1177/014920630102700202](https://doi.org/10.1177/014920630102700202)
- Weick KE, Sutcliffe KM (2001) *Managing the unexpected*. Jossey-Bass, San Francisco
- Williams KY, O'Reilly CA (1998) Demography and diversity in organizations: a review of 40 years of research. In: Staw B, Sutton R (eds) *Research in organizational behavior*, vol 20. JAI Press, Greenwich, pp 77–140
- Winquist J, Larson J (1998) Information pooling: when it impacts group decision making. *J Pers Soc Psychol* 74:371–377. doi:[10.1037/0022-3514.74.2.371](https://doi.org/10.1037/0022-3514.74.2.371)
- Wittenbaum G, Stasser G (1996) Management of information in small groups. In: Nye J, Brower A (eds) *What's social about social cognition? Social cognition research in small groups*. Sage, Thousand Oaks, pp 3–28

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.