Cognition 142 (2015) 191–204

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

The kind of group you want to belong to: Effects of group structure on group accuracy

Martin L. Jönsson [a,*], Ulrike Hahn [b], Erik J. Olsson [a]

[a] Department of Philosophy, Lund University, Sweden
[b] Department of Psychological Sciences, Birkbeck University of London, United Kingdom

ARTICLE INFO

Article history:
Received 18 December 2014
Revised 11 April 2015
Accepted 18 April 2015
Available online 2 June 2015

Keywords:
Wisdom of crowds
Group structure
Group accuracy
Information flow
Social networks
Group learning

ABSTRACT

There has been much interest in group judgment and the so-called 'wisdom of crowds'. In many real world contexts, members of groups not only share a dependence on external sources of information, but they also communicate with one another, thus introducing correlations among their responses that can diminish collective accuracy. This has long been known, but it has—to date—not been examined to what extent different kinds of communication networks may give rise to systematically different effects on accuracy. We argue that equations that relate group accuracy, individual accuracy, and group diversity (see Hogarth, 1978; Page, 2007) are useful theoretical tools for understanding group performance in the context of research on group structure. In particular, these equations may serve to identify the kind of group structures that improve individual accuracy without thereby excessively diminishing diversity so that the net positive effect is an improvement even on the level of collective accuracy. Two experiments are reported where two structures (the complete network and a small world network) are investigated from this perspective. It is demonstrated that the more constrained network (the small world network) outperforms the network with a free flow of information.

1. Introduction

Interaction with others in different social groups is an essential part of the human condition. As members of a jury we deliberate with fellow jurors in order to arrive at an appropriate verdict, as members of a legislative body we interact with others to create and repeal laws, and as members of research groups we pool our resources so that we jointly can perform better than we can do individually.

We frequently trust the verdicts and estimates of our groups, even in cases where they are in conflict with our own. The well-foundedness of this trust has been the subject of much research in social psychology. Early on, Galton (1907) famously compared the accuracy of a group with that of its members in guessing the weight of an ox during a stock and poultry exhibition. During subsequent decades, social psychologists carried on in this tradition by comparing groups with their members on a variety of tasks, from the estimation of room temperature, to the judgment of children's intelligence from photographs, to the solution of mathematical problems (see, e.g. Knight, 1921; Shaw, 1932, for extensive reviews see Gigone & Hastie, 1997; Hill, 1982; Lorge & Brenner, 1958). The bottom line of much of this was that, on the one hand, results could not really be made sense of without formal statistical tools, and, on the other, that once these were properly utilized, much of these earlier results seemed trivial.

In the words of Gigone and Hastie (1997).

> Statistical combinations of judgments have long been known to cancel out unsystematic judgment error

* Corresponding author.

http://dx.doi.org/10.1016/j.cognition.2015.04.013
0010-0277/

(Hogarth, 1977). The standard error of the mean of several judgments is smaller than the standard deviation of the judgments themselves; groups almost inevitably outperform their members simply by averaging those members' judgments. Such accuracy gains can hardly be attributed to anything special about the group judgment process; the group need not meet at all.

Gigone and Hastie (1997: 159)

So the real question of group research must be the extent to which the group is better than the statistical aggregate (Gigone & Hastie, 1997). Or, to put this differently, what is it that the group adds?

One way of approaching this question is to manipulate communication channels within a group and examine attendant effects. Experimental manipulation of the information participants receive from others allows inference about the extent to which they use that information. It thus provides a methodological window into how people go about combining what they believe with information they receive from others.

This question seems at least as relevant now as it did in the early days of small group research, because it has become ever more apparent that our beliefs and opinions are determined not merely by our own observations, but, to an arguably even greater extent, by the evidence we receive through the testimony of others (see e.g., Coady, 1992). Consequently, there is only so much one can study about human learning, judgment and decision making without taking into account the social dimension of belief formation (see also, Goldstone & Gureckis, 2009).

This in turn suggests a subtle shift in emphasis concerning the kinds of groups and tasks that are of interest and what aspects of group influence and performance seem most worthy of examination. Much of the past research on groups (as surveyed in the reviews of Gigone & Hastie, 1997; Hill, 1982; Lorge & Brenner, 1958) has focussed on the quality of the group response itself, and this is also the central theme in the recent revival of this tradition of research under the header of 'wisdom of crowds' (Hertwig, 2012; Herzog & Hertwig, 2009; Surowiecki, 2004). However, it is at least as interesting and important to ask what the group does for the individual, and how this develops, that is, to ask not just how group performance compares to individual performance but to ask how both individual and group performance are changed by group communication.

It is here that useful links can be formed with the burgeoning literature on networks, in particular social networks (for an introduction see e.g., Jackson, 2010). Patterns of communication between individuals in groups give rise to network structure (see also, Goldstone, Roberts, & Gureckis, 2008; McGrath, Arrow, & Berdahl, 2000): depending on context, all members may be exchanging views and listening to one another freely; alternatively, only some members may be communicating directly with one another. Finally, even where all individuals hear all information being exchanged, selective attention and weighting (see e.g., Friedkin & Johnsen, 1999) of others' information (determined, for example, by perceived competence) imposes an effective network structure to the communication that diverges from the surface level whereby everyone is communicating with everybody else.

As just indicated, experimental manipulation of the structure of communication may provide insight into what it is that being part of a group is adding. At the same time, it raises interesting questions of its own concerning the extent to which different types of communication networks may systematically differ in their impact on our beliefs (on the general benefits of taking a network perspective to traditional group research see also, Katz, Lazer, Arrow, & Contractor, 2004).

To this end, we present two experimental studies manipulating the communication structure within a group and examining its impact on the accuracy of participants' beliefs. To sidestep some of the pitfalls of the early work on group accuracy, our analysis is informed by two equations that relate group validity, individual validity and group diversity. These equations demonstrate—for two different ways of aggregating opinion and two different ways of understanding accuracy—the conditions under which the group will outperform its average individual member by mathematical necessity. First, work by Ghiselli (1964, chap. 7) and Hogarth (1978) points out that if the validity of a sequence of estimates is understood in terms of the correlation between it and the true values, the validity of the group estimate can be shown to always exceed the average validity of the answers of the group members as long as the members are not perfectly correlated with each other and error is unbiased.

More precisely, if we let n be the number of group members, $s_{x_i}, s_{\bar{x}}, s_t$ be sequences of the estimates of group member i, mean estimates, and true values respectively, and $\rho_{x,y}$ be the correlation between two sequences x and y, Hogarth's equation states that[1]

$$\rho_{s_t s_{\bar{x}}} = \frac{\sqrt{n}\frac{\sum_{i=1}^{n}\rho_{s_t s_{x_i}}}{n}}{\sqrt{1+(n-1)\frac{\sum_{i=1}^{n}\sum_{j=i+1}^{n}\rho_{s_{x_i} s_{x_j}}}{n}}} \quad (1)$$

The limiting case where $n = \infty$ is captured by following equation:

[1] It should be noted that understanding validity in terms of a correlation results in a fairly coarse-grained concept of validity. For instance, assume that Bob and Sue have answered in the following way:

	Bob	Sue	Correct
Question 1	13	5	5
Question 2	15	7	7
Question 3	11	3	3
Question 4	13	5	5

On the correlational understanding of validity, Bob's and Sue's answers are, counterintuitively, equally valid (both answers are perfectly correlated with the correct answer). This might be what Hogarth is after when he remarks that his results only hold in circumstances where 'the judgmental task consists of rank ordering alternatives—that is the *level* of judgment is not important.' (Hogarth, 1978: 41, emphasis in original). Nonetheless, even when the exact values are important for a correct answer, the correlation between a sequence of answers and the correct answers gives us an indication of how good the answers are; answers that are very poorly correlated with the correct answers cannot be correct.