



# Optimizing group learning: An evolutionary computing approach

Igor Douven

## ► To cite this version:

Igor Douven. Optimizing group learning: An evolutionary computing approach. Artificial Intelligence, 2019, 275, pp.235-251. 10.1016/j.artint.2019.06.002 . hal-02964488

**HAL Id: hal-02964488**

**<https://hal.sorbonne-universite.fr/hal-02964488>**

Submitted on 12 Oct 2020

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/333882477>

# Optimizing Group Learning: An Evolutionary Computing Approach

Article in *Artificial Intelligence* · June 2019

DOI: 10.1016/j.artint.2019.06.002

CITATIONS

3

READS

185

1 author:



Igor Douven

French National Centre for Scientific Research

197 PUBLICATIONS 2,364 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Explanatory Reasoning [View project](#)



Experimental Philosophy [View project](#)

# Optimizing Group Learning: An Evolutionary Computing Approach

Igor Douven  
SND / CNRS / Sorbonne University

## Abstract

We study groups of interacting agents who are trying to discover probabilistic truths on the basis of sequentially provided evidence and information about the belief states of other group members. The main research question is which combination of epistemic principles—combinations of an evidential update rule, a rule for determining peerhood, and a rule for aggregating probability functions—such groups should adopt to strike the best balance between being fast and being accurate, where the former is understood by reference to how long it takes before a majority of the group assigns a high probability to the true hypothesis, and the latter by reference to the average Brier penalty incurred by the group. The main methodology to be used is that of agent-based optimization, which is a specific form of evolutionary computing. We implement this methodology in a generalization of the Hegselmann–Krause model. In the end, we are able to identify optimal procedures for taking into account both direct evidence and information about one’s peers’ beliefs. At the same time, we note that optimality for such procedures is dependent on context.

Keywords: agent-based modeling; computer simulations; explanatory reasoning; evolutionary computing; group learning; Hegselmann–Krause model; opinion dynamics.

## 1 Introduction

There are distinctively social aspects to learning. Not only do we learn *from* others—as when we learn from our parents or our teachers—we also often can acquire new knowledge only if we collaborate *with* others, for instance, as members of the same laboratory or research group (Goldman, 1999). Indeed, modern science is largely a community effort; scientists who might want to work in relative isolation—in the manner of Cavendish, Darwin, or Newton—would nowadays stand little chance of making any appreciable progress (Gribbin, 2002, p. 359).

Once the communal aspects of knowledge acquisition are recognized, it becomes natural to ask how we might best organize a group of agents bent on a joint research goal, supposing we want to optimize the group’s performance. This is a broad question, of course, which could concern topics ranging from after-work socializing to decisions about how to allocate grant money or other resources within a group (Kitcher, 1993). The present paper, instead, focuses on some of the most fundamental *epistemic* mechanisms that groups might decide to put in place, or might be encouraged to adopt, for how their members ought to act and especially to interact. In particular, it focuses on mechanisms for how group members should respond to, or “update on,” new evidence while also being informed about the changing belief states of others in the group, who they know to acquire evidence pertaining to a common source (all group members study the same economic phenomenon, say, or the same virus, or the same genetic mutation).

The notion of *optimizing* group performance will also be understood in a strictly epistemic sense, that is, as relating to the question of how to get at the *truth* of whatever issue the group is working on. While truth is generally regarded as the overarching goal of science (Psillos, 1999), it is equally acknowledged that science serves many practical purposes, too. For that reason, we are often forced to make a speed–accuracy trade-off: we do want to get at the truth, but we also want to get there fast, which may require that we settle on quickly becoming *highly confident* in the truth (even if we cannot completely rule out alternative hypotheses) or quickly getting *close* to the truth (Kuipers, 2000, 2001, 2014, 2019).

Thus, the question to be investigated is how members of a research group should update on the receipt of new evidence in a social setting, where they also have access to relevant beliefs of (some or all of) their colleagues, supposing the group wants to strike the best balance between speed (getting at the truth fast) and accuracy (minimizing error rates). The main methodological tool to be used is that of computational agent-based modelling, which has become a central topic in artificial intelligence (Shoham, Powers, & Grenager, 2007; Tamargo, Garcia, Falappa, & Simari, 2014; Nunes & Antunes, 2015; Gottifredi et al., 2018). Specifically, we build on the well-known Hegselmann–Krause (HK) model for studying opinion dynamics in groups of interacting agents focused on a common research question (Hegselmann & Krause, 2002, 2005, 2006, 2009; for related models, see Deffuant et al., 2000; Dittmer, 2001; Weisbuch et al., 2002; Pluchino, Latora, & Rapisarda, 2006; Semeshenko, Gordon, & Nadal, 2008; De Langhe & Greiff, 2010). In that model, each agent holds only a single opinion at any point in time. Later work has generalized the model to ones where agents could have multiple opinions (e.g., Lorenz, 2003, 2008; Riegler & Douven, 2009; Wenmackers, Vanpoucke, & Douven, 2012, 2014). Recently, Douven and Wenmackers (2017) have presented an extension of the HK model in which each agent, at any moment, assigns probabilities to all members of a set of relevant possibilities.

Douven and Wenmackers’ (2017) aim was to compare different update rules (rules for adapting probabilities in response to new evidence) within a social setting. Specifically, they compared Bayes’ rule (Oaksford & Chater, 2007) with an update rule intended to formalize the kind of explanatory reasoning that has lately been much in the limelight both in artificial intelligence (Konolige, 1992; Boutilier & Becher, 1995; Baral, 2000; Lin & You, 2002; Glass, 2018a, 2018b; Teijeiro & Félix, 2018) and in cognitive psychology (Koslowski et al., 2008; Bes et al., 2012; Williams & Lombrozo, 2013; Legare & Lombrozo, 2014; Lombrozo & Gwynne, 2014; Douven & Schupbach, 2015a, 2015b; Lombrozo, 2016; Johnston et al., 2017; Douven & Mirabile, 2018; Koslowski, 2018). One way in which the present paper goes beyond Douven and Wenmackers’ work is by considering a number of different formalizations of explanatory reasoning, particularly ones that were motivated by the aforementioned recent work in cognitive psychology. Another generalization concerns the way an agent determines which of her group members she takes to be her *epistemic peers* (Douven, 2010; De Langhe, 2013), which in Douven and Wenmackers (2017) are the members whose probabilities she takes into account in computing her new probabilities. While Douven and Wenmackers used one particular metric for defining epistemic peerhood, we consider others as well. Furthermore, we compare different ways in which group members may accommodate the information about their peers’ probabilities, thereby also going beyond previous research, Douven and Wenmackers having considered only the linear pooling of probabilities (Stone, 1961; for details, see below).

Where Douven and Wenmackers’ main conclusion was that, at least in a social setting, explanatory reasoning can be superior to Bayesian updating both in terms of speed and in terms of accuracy, the present research rather aims to show how we may be able to find an optimal combination of update rule, rule for determining peerhood, and rule for taking into account one’s peers’ probabilities, where optimality is still a matter of balancing speed and accuracy. The main method to be used is evolutionary

computing (Holland, 1975; Koza, 1992), in which combinations of the aforementioned kind compete for survival and reproduction, their success depending on how fast and accurate they are.

It is important to be clear about the motivation for the current research. Recently, much has been made of the so-called wisdom-of-crowds effect, the often-reproduced finding that averaged group estimates of some parameter value tend to be closer to the truth not only than the individual estimates but even than expert estimates (Surowiecki, 2004; Page, 2007). At the same time, Lorenz et al. (2011) report evidence showing that if group members are made aware of each others' estimates, then that will not only lead members to adjust their own estimates, the various resulting adjustments also will tend to have a negative effect on the accuracy of the averaged group estimate. This might seem to give a reason for trying to *prevent* the kind of interactions among group members that were experimentally explored by Lorenz and coauthors. In a similar vein, Zollman (2007) found that while *increasing* the connectivity in an epistemic network tends to *increase* the speed with which the network converges on an opinion, it at the same time tends to *decrease* the chances that the network will arrive at the truth (see also Kummerfeld & Zollman, 2016, and Hahn, Hansen, & Olsson, 2019; but cf. Mason & Watts, 2012, and Rosenstock, Bruner, & O'Connor, 2017).

However, while these authors found a potential negative effect of epistemic interactions among the members of research groups, it is to be noticed that such interactions are not *bound* to be detrimental. To the contrary, as Douven and Wenmackers' (2017) results made manifest, whether the influence will have a positive or a negative effect on the accuracy of the group estimate may all depend on *how* the individual group members' beliefs are influenced by those of their peers. There are update procedures which take into account both new information about the parameter value and information about the beliefs about that value held by others in the group and which greatly increase not only the accuracy of the group estimate but also the speed with which that estimate converges to the true value, supposing group members are repeatedly offered new evidence as well as new information about the beliefs of their peers. Groups might decide to adopt those update procedures to counteract the kind of effects of social influence discovered by Lorenz et al. (2011) and others. Naturally, that presupposes that such procedures have been identified. The present paper describes in some detail how this could be accomplished.

Section 2 outlines the theoretical background of the present research, in particular, empirical and formal work on update rules, and computational work on opinion dynamics. Section 3 starts by summarizing the opinion dynamics framework used in Douven and Wenmackers (2017) and proposes a generalization of this framework. Section 4, finally, explores the generalized framework and uses it to identify optimal procedures for taking into account both direct evidence and information about one's peers' beliefs.

## 2 Theoretical background

We are going to look at groups of agents, each of whom repeatedly updates her belief state upon the receipt of new information. The new information will be of two kinds: on the one hand, agents receive new information directly about the subject matter of their research, while, on the other hand, they receive new information about how others in their group changed their belief states after these others received new "direct" information about the same subject matter. All agents will always first receive a piece of direct information—which may be different for different agents, even though the information will come from the same source—and then, after each of them has updated on that, they will receive information about the new belief states of their colleagues, which may occasion a further adjustment of their own belief state. Both parts of this process rely on procedures that have been widely discussed

in the literature, the first part involving so-called update rules, the second involving a variant of the already mentioned HK model.

## 2.1 Update rules

It is customary to distinguish categorical (yes/no) beliefs from graded beliefs. While there is interesting work on how we change the former kind of beliefs when we obtain new information (see, e.g., Alchourrón, Gärdenfors, & Makinson, 1985), there is widespread agreement that, at least in scientific contexts, the more important belief type is the graded one. Graded beliefs are usually conceived as subjective probabilities (Jeffrey, 1965; Howson & Urbach, 1989; Earman, 1992; Oaksford & Chater, 2007), though they can also reflect objective chances (Lewis, 1980). The agents who constitute the communities considered in this paper will all have, at any point in time, such graded beliefs; no categorical beliefs will be attributed to them. Accordingly, the only update rules that will matter for our purposes are ones that operate on probabilities.

Among these rules, Bayes' rule has been attracting most of the attention for a while now. According to this rule, upon receiving new information we ought to make our previous *conditional* probabilities on the supposition of that information into our new *unconditional* probabilities. More formally, where  $\mathcal{A}$  is an algebra of propositions,  $\phi \in \mathcal{A}$  the strongest proposition an agent learns at a given point in time, and  $\Pr(\cdot)$  and  $\Pr_\phi(\cdot)$  the agent's probability assignments to the propositions in  $\mathcal{A}$  immediately before and immediately after learning  $\phi$ , respectively, it ought to hold for all  $\psi \in \mathcal{A}$  that  $\Pr_\phi(\psi) = \Pr(\psi | \phi)$ , provided  $\Pr(\phi) > 0$ , with  $\Pr(\psi | \phi)$  defined as  $\Pr(\phi \wedge \psi) / \Pr(\phi)$ .

Bayes' rule is simple yet extremely powerful, as witnessed by the remarkable successes of Bayesian statistics (Gelman et al., 2013; Kruschke, 2015) and of various applications of the Bayesian machinery in artificial intelligence (e.g., Pearl, 2000; Kern-Isberner & Lukasiewicz, 2004; Korb & Nicholson, 2004; Darwiche, 2009). While that justifies the popularity of Bayes' rule, researchers have recently started to study some similar-looking rules which, however, attribute special weight to *explanatory* relationships between what one learns and possible answers to one's questions of interest. These other update rules can be conceived as formalizations of an old philosophical idea, which commonly goes by the name of "Inference to the Best Explanation" (IBE). According to this idea, explanatory considerations can give grounds for belief; or in terms of graded belief, a hypothesis being the best explanation of one's evidence licenses increased confidence in that hypothesis.

Among philosophers of science, IBE was once widely heralded as a cornerstone of scientific methodology (Boyd, 1984, 1990; McMullin, 1992). But interest in IBE faded with the rise of Bayesianism in philosophy. This was especially so when philosophers came to endorse (rather uncritically—see below) arguments to the effect that Bayes' rule is the only rational update rule, so that, in particular, changing one's belief state on the basis of explanatory considerations, as IBE would lead one to do, has to be irrational, unless the belief change is somehow reducible to a Bayesian update, in which case IBE is redundant (van Fraassen, 1989). Reasons for the renewed interest that researchers (apart from philosophers now also psychologists and computer scientists) have taken in IBE are of a normative and a descriptive variety.

As for the former, the main arguments given in favor of Bayes' rule being the only rational update rule are the so-called dynamic Dutch book argument (Teller, 1973) and the inaccuracy-minimization argument (Rosenkrantz, 1992). At bottom, both arguments purport to show that updating in some non-Bayesian fashion comes at a potential cost, monetary in the case of the Dutch book argument (inflicted via a collection of bets that individually seem fair to the non-Bayesian updater but that jointly guarantee a negative net payoff) and epistemic in the case of the inaccuracy-minimization argument (the claim being that the expected inaccuracy of the non-Bayesian's updated subjective probabilities is

greater than it would have been were she a Bayesian). What various authors began to realize, however, is that, even granting these arguments, Bayesians had failed to show that the costs were incurred *for no benefit*—that that was so had just been tacitly assumed. Instead, these authors showed that, under the right circumstances, a variety of non-Bayesian update rules, notably ones meant to formalize IBE, outperform Bayes’ rule in terms of speed of convergence to the truth (how fast does the true hypothesis receive a high enough probability, in some sense of “high enough”?), in terms of accuracy (as operationalized in terms of some scoring rule; see below), or even in both respects (Glass, 2007, 2012, 2018a, 2018b; Douven, 2013, 2017; Douven & Wenmackers, 2017; Trpin & Pellert, 2019).

The rules that will be relevant to our present research are, next to Bayes’ rule, precisely such forms of explanatory reasoning. While these rules give credit to hypotheses for their explanatory power, they are otherwise very similar to Bayes’ rule; in particular, they also yield probabilities as output when given probabilities as input. Specifically, we will compare with Bayes’ rule and with each other three rules: a rule proposed in van Fraassen (1989) and dubbed “EXPL” in Douven (1999), and two variants of that rule dubbed “Good’s rule” and “Popper’s rule” in Douven (2017).

Where  $\Psi = \{\psi_i\}_{1 \leq i \leq n}$  is a set of mutually exclusive and jointly exhaustive self-consistent hypotheses,  $\Pr(\cdot)$  and  $\Pr_\phi(\cdot)$  are as defined before, and  $\Pr(\phi) > 0$ , EXPL is the rule according to which

$$\Pr_\phi(\psi_i) = \frac{\Pr(\psi_i) \Pr(\phi | \psi_i) + \mathcal{E}(\psi_i, \phi)}{\sum_{j=1}^n (\Pr(\psi_j) \Pr(\phi | \psi_j) + \mathcal{E}(\psi_j, \phi))}.$$

Here,  $\mathcal{E}(\cdot, \cdot)$  assigns a bonus  $c (> 0)$  to the member of  $\Psi$  that explains the evidence  $\phi$  best, and nothing to the other hypotheses. Notice that, as stated here, EXPL is actually a schema, yielding a different rule for each  $c$ ; also notice that if  $c$  were set to 0, EXPL would coincide with Bayes’ rule. In this paper, we follow van Fraassen in understanding the notion of explanatory bestness as it occurs in EXPL in terms of likelihood maximization, in that  $\psi_i \in \Psi$  is said to be the best explanation of  $\phi$  iff  $\Pr(\phi | \psi_i) > \Pr(\phi | \psi_j)$  for all  $\psi_j \in \Psi$  with  $j \neq i$ . (If two or more hypotheses are tied for best explanation, the bonus  $c$  will be split evenly among them.)

In EXPL, the best explanation gets all the credit while the other hypotheses get nothing, even if they are perhaps good explanations of the evidence, just not the best. Good’s rule and Popper’s rule, now to be stated, are different in that they look at the explanatory power of each single hypothesis and attribute some bonus (or possibly also malus) *proportional* to that power, so that the best explanation will get the largest bonus—with how large it is depending on how good an explanation it is—but good-but-non-best explanations will still receive *some* bonus, and *bad* explanations may get a malus.

The rules are instances of the schema

$$\Pr_\phi(\psi_i) = \frac{\Pr(\psi_i) \Pr(\phi | \psi_i) + c \Pr(\psi_i) \Pr(\phi | \psi_i) \mathcal{M}(\psi_i, \phi)}{\sum_{j=1}^n (\Pr(\psi_j) \Pr(\phi | \psi_j) + c \Pr(\psi_j) \Pr(\phi | \psi_j) \mathcal{M}(\psi_j, \phi))}.$$

Good’s rule is obtained from this by substituting for  $\mathcal{M}(\cdot)$  the following measure, taken from Good (1960), according to which hypothesis  $\psi$ ’s power to explain evidence  $\phi$  is given by

$$\ln \left( \frac{\Pr(\phi | \psi)}{\Pr(\phi)} \right),$$

while substituting for  $\mathcal{M}(\cdot)$  Popper’s (1959) measure of the explanatory goodness of hypothesis  $\psi$  in light of evidence  $\phi$ ,

$$\frac{\Pr(\phi | \psi) - \Pr(\phi)}{\Pr(\phi | \psi) + \Pr(\phi)},$$

yields Popper’s rule.

Thus defined, Good's rule and Popper's rule are also only schemata; to obtain a specific instance of either, we must fix a value for  $c$  ( $> 0$ ), which then determines, for all  $i$ , the percentage of  $\psi_i$ 's probability that, after a Bayesian update on  $\phi$ , is added, in proportion to  $\psi_i$ 's power to explain  $\phi$ . Also note that we would again obtain Bayes' rule by setting  $c = 0$  in Good's or Popper's rule. (For further details, see Douven, 2017.)

A growing awareness that there is something amiss with the main arguments in favor of Bayes' rule—backed by evidence that rules like EXPL, Good's rule, and Popper's rule can, under certain circumstances, outperform Bayes' rule—gave normative grounds for taking explanatory reasoning more seriously than many had long been willing to do. As said, however, there were also descriptive reasons for the renewed interest.

While there is clear evidence showing that people do, in some types of situations, comply to Bayesian principles, at least approximately (Oaksford & Chater, 2007; Tenenbaum et al., 2011; Griffiths, Lieder, & Goodman, 2015), there is equally clear evidence that, in other types of situations, people tend to deviate starkly from those same principles (Phillips & Edwards, 1966; Edwards, 1968; Marks & Clarkson, 1972; Fischhoff & Lichtenstein, 1978; Schum & Martin, 1982). Such deviations could be entirely unsystematic,<sup>1</sup> but a number of them have turned out to be due to people giving weight, in a non-Bayesian fashion, to explanatory connections between evidence received and whichever hypotheses were of interest to them. See, for instance, Pennington and Hastie (1988, 1992, 1993) and Bes et al. (2012). Still more directly relevant to the current project, Douven and Schupbach (2015a) found in a number of experiments that people's judgments of explanatory goodness were strongly predictive of their subjective probability judgments, much more so than the objective probabilities pertaining to the situation, which, according to Bayesians, should have been the only significant predictor in this case. In a re-analysis of their data, Douven and Schupbach (2015b) even found that explanatory considerations had influenced their participants' belief change more or less as if they relied on a measure of explanatory power very similar to Good's or Popper's.

## 2.2 The HK model and beyond

Computational agent-based modeling has become a mainstay in the social sciences, used to study a great variety of phenomena and processes—the emergence of social norms, racial segregation, voting behavior, and so on—that are too complex to be investigated using strictly analytical tools. The technique is also rapidly gaining popularity in the study of social aspects of knowledge and belief, fueled by the observation already hinted at above, to wit, that our being socially embedded in one or more communities of truth-seekers is crucial to our epistemic success.

Arguably, the most widely used computational agent-based model in this field is the Hegselmann–Krause (HK) model. Hegselmann and Krause themselves mainly used the model to address *descriptive* issues, such as the conditions under which initially diverging opinions in a community of interacting agents tend to converge or when they tend to polarize. More recently, the model has also been applied to a variety of *normative* issues, such as questions concerning reliabilism and truthful assertion (Olsson, 2008), peer disagreement (De Langhe, 2013), and efficient truth approximation in science (Douven & Kelp, 2011).

In the version of the HK model that is most of interest to us (though not in the most basic, original version of the model), communities of agents seek to determine the value  $\tau$  of some unspecified

---

<sup>1</sup>Or they could have been systematic but evidence of general cognitive biases or defects. For instance, Costello and Watts (2016, 2018) have tried to explain away some deviations from Bayesian norms as being due to sampling errors, which in turn they attribute to human memory limitations. However, this strategy failed in the face of some deviations from Bayesian updating; see Douven (2019).



parameter, where it is given that  $\tau \in (0, 1)$ . The agents update their estimates of  $\tau$  repeatedly, at discrete time steps, where each agent updates on the basis of (i) information she receives about  $\tau$ , and (ii) the estimates of  $\tau$  of those agents who are within her bounded confidence interval (BCI), meaning that their estimate is within some distance  $\varepsilon$  of the agent's own estimate. More formally, the estimate of agent  $x_i$  after the  $(u + 1)$ -st update is given by

$$x_i(u + 1) = \alpha \frac{1}{|X_i(u)|} \sum_{j \in X_i(u)} x_j(u) + (1 - \alpha)\tau,$$

with  $x_j(u)$  the estimate of agent  $x_j$  after the  $u$ -th update,

$$X_i(u) := \{j : |x_i(u) - x_j(u)| \leq \varepsilon\}$$

the set of agents within agent  $x_i$ 's BCI after the  $u$ -th update, and  $\alpha \in [0, 1]$  a parameter determining the weight the agent gives to the “social” part of the updating relative to the “evidential” part.<sup>2</sup>

One reason why the HK model is so popular is its great flexibility. In general, it takes little effort to extend the model or to adapt it to a researcher's specific needs. For example, Douven (2010), Douven and Riegler (2010), and De Langhe (2013) use simple extensions of the model to study situations in which agents receive evidence beset by some degree of measurement error. Crosscombe and Lawry (2016) present another extension to study agents with (partially) vague beliefs. For still other extensions, see Lorenz (2003, 2008), Jacobmeier (2004), Pluchino, Latora, and Rapisarda (2006), Riegler and Douven (2009), and Wenmackers, Vanpoucke, and Douven (2012, 2014).

We will be mostly interested in an extension of the HK model proposed in Douven and Wenmackers (2017). In this extension, agents' belief states at a given time are not single real numbers (as in the original model) but probability functions, defined on a set of self-consistent, mutually exclusive, and jointly exhaustive hypotheses  $\Psi = \{\psi_i\}_{1 \leq i \leq n}$ . It is still the case, though, that the agents take into account both evidence directly pertaining to the members of  $\Psi$  and the belief states (i.e., probability functions) of those in their BCI.

Naturally, in this model the BCI has a definition different from the one stated above. Here, two agents are in each other's BCI precisely if the sum of the absolute differences in the probabilities they assign to the hypotheses is below a given threshold value, or formally, agent  $i$  with belief state  $\text{Pr}_i(\cdot)$  is within the BCI of agent  $j$  with belief state  $\text{Pr}_j(\cdot)$  precisely if  $\Delta(\text{Pr}_i(\cdot), \text{Pr}_j(\cdot)) \leq \varepsilon$ , where

$$\Delta(\text{Pr}_i(\cdot), \text{Pr}_j(\cdot)) := \sum_{k=1}^n |\text{Pr}_i(\psi_k) - \text{Pr}_j(\psi_k)|.$$

Note that, given this definition, an agent is always in her own BCI (as is also the case in the original HK model and in the extensions mentioned above).

The actual updating in Douven and Wenmackers' model then proceeds as follows: In a first, “evidential,” step, agent  $j$  updates her probabilities on a new piece of evidence she receives, using either Bayes' rule or EXPL; that is, where  $\text{Pr}_j^u(\cdot)$  is her probability function after the  $u$ -th update, she first forms  $\text{Pr}_j^{\text{evid}}(\cdot)$  via either Bayes' rule or EXPL. Then, in a second, “averaging,” step, the agent averages over the probability functions of the agents in her resulting BCI,  $X(\text{Pr}_j^{\text{evid}}(\cdot))$ , which yields  $\text{Pr}_j^{\text{av}}(\cdot)$ . Finally, she sets  $\text{Pr}_j^{u+1}(\cdot) = \text{Pr}_j^{\text{av}}(\cdot)$ .

The averaging takes place by what is known as “linear pooling,” an idea formally introduced in Stone (1961), meaning that the agents take *straight arithmetic averages* of the probabilities of the agents

<sup>2</sup>In this model, and also in Douven and Wenmackers' extension of it to be discussed below, all agents in the BCI are treated on a par. If information about the credibility of these agents were available, one might plausibly wish to weigh them differently, for instance, in the manner of Tamargo et al. (2014) and Gottifredi et al. (2018).

within their BCI. Formally, for all  $\psi \in \Psi$  and all  $\Pr_k^{\mu_{\text{evid}}}(\cdot) \in X(\Pr_j^{\mu_{\text{evid}}}(\cdot))$ , it holds that

$$\Pr_j^{\mu_{\text{av}}}(\psi) = \frac{\sum_k \Pr_k^{\mu_{\text{evid}}}(\psi)}{|X(\Pr_j^{\mu_{\text{evid}}}(\cdot))|}.$$

For the reasons explained in Douven and Wenmackers, the resulting function again satisfies the axioms of probability theory.

Douven and Wenmackers' (2017) explicit aim was to compare Bayes' rule with EXPL in a social setting. Previous research using computer simulations had shown that, in an individual setting (and so without any averaging of probability functions taking place), EXPL tends to outperform Bayes' rule in terms of speed, but Bayes' rule tends to outperform EXPL in terms of accuracy (Douven, 2013). That would seem to make it a matter of context which of these rules we ought to use (supposing the choice is between just them), given that we sometimes seem to care more about speed, and sometimes more about accuracy.

But Douven and Wenmackers wondered whether a speed–accuracy trade-off would also have to be made if agents were to change their probabilities not just on the basis of direct evidence but also as a result of social influence, as in reality people are known to do (e.g., Miller & McFarland, 1991; Mason, Conrey, & Smith, 2007; Lorenz et al., 2011). To address this question, they ran computer simulations using the extension of the HK model just described, with groups of simultaneously and repeatedly updating agents who received evidence that was, while coming from the same source (i.e., distribution) for all agents, randomly generated per agent. Douven and Wenmackers systematically studied different communities of agents, where the communities differed along two dimensions, to wit, by the value for  $c$  they substituted in EXPL (with those setting  $c$  equal to 0 *de facto* amounting to Bayesians), and by the value of  $\varepsilon$ , that is, by their being more or less liberal in considering others as their epistemic peers.

The agents in Douven and Wenmackers' simulations all receive binomial data (i.e., independent and identically distributed binary data), the data points coming in one at a time, and their goal is to estimate on that basis the probability of “success.” For the sake of concreteness, one may imagine the data to consist of the outcomes of repeated tosses of a coin with unknown bias, where this bias (say, the probability of heads) can take eleven possible values, ranging from 0 (certain to land tails) to 1 (certain to land heads), in increments of .1. Initially, all these values are equally likely to indicate the true bias of the coin. Thus, at the start of the updating process, it is reasonable for the agents to invest equal confidence in all given bias possibilities. Importantly, while each agent in the community receives data from a different coin, it is known to them that all coins have the *same* bias. The community of agents is tasked to estimate this bias.

As said, the relevant dimensions of evaluation were speed and accuracy. The former was operationalized in terms of the percentage of the community that assigns a high probability to the truth (where “high probability” was understood as a probability above .9, but other plausible values would have done as well), and the latter in terms of the so-called Brier score (Brier, 1950), which is the most commonly used scoring rule. Given again a set  $\{\psi_i\}_{1 \leq i \leq n}$  of mutually exclusive and jointly exhaustive self-consistent hypotheses, with  $\llbracket \psi_i \rrbracket \in \{0, 1\}$  indicating the truth value of  $\psi_i$  (0 = false, 1 = true), an agent assigning probability  $\Pr(\psi_i)$  to  $\psi_i$  (for each  $i \leq n$ ) incurs a Brier penalty of  $\sum_{i=1}^n (\llbracket \psi_i \rrbracket - \Pr(\psi_i))^2$ .

In their simulations, Douven and Wenmackers effectively compared 121 (= 11 × 11) update procedures with each other, one for each combination of  $c$  and  $\varepsilon$ , with both being in  $\{0, 0.1, 0.2, \dots, 1.0\}$ . The communities these authors considered all consisted of 50 agents, each of whom received, per simulation, the outcomes from 500 tosses. Given that, as previous research had shown, results may depend on the




Figure 1: Average Brier scores for six bias possibilities ( $p = x$  indicates that the chance for each of the coins to come up heads equals  $x$ ).

specific bias of the coin, simulations were run for biases  $p \in \{0, .1, .2, .3, .4, .5\}$ .<sup>3</sup> For each of these bias possibilities, 50 simulations were run.

In all simulations, average Brier scores per agent were calculated after each update. Per simulation, these averages were then summed over all 500 tosses, and finally the average over the 50 simulations was calculated. The results are graphically summarized in Figure 1, which is reproduced from Douven and Wenmackers (2017). The bottom row of each of the six plots represents the communities of Bayesian updaters, where these only differ by their willingness to count others as their epistemic peers; the square in the bottom-left corner of the plots represents communities of “pure” Bayesians, who attach no value to explanatory strength and are unwilling to be influenced by anyone in their group. As Douven and Wenmackers observe, among the Bayesian communities the pure ones do best on the count of Brier scores: as is clearly visible, Bayesians who *are* willing to be influenced by others incur higher Brier penalties, on average. More significantly, *non*-Bayesian updaters who are willing to be influenced by others do better still than pure Bayesians. (For details, see Douven & Wenmackers, 2017.)

As for speed, it had already been found that, on average, and in the statistical model considered in the simulations, EXPL leads agents to assign a high probability to the truth faster than does Bayes’ rule. This result was reproduced in the kind of social setting assumed by Douven and Wenmackers. The plots in Figure 2 show graphically for different combinations of  $c$  and  $\epsilon$  the average percentages of agents (averaged over the simulations) that assign a high probability to the truth, for the relevant bias possibilities, and after each toss (compare Figure 6 in Douven & Wenmackers, 2017).

Douven and Wenmackers (2017) were interested in comparing Bayes’ rule and EXPL in a social setting, along the dimensions of speed and accuracy. They found that, while in an individual setting, the two rules appeared to have different strengths and weaknesses—Bayes’ rule being more accurate

<sup>3</sup>For reasons of symmetry, it would have been superfluous to also run simulations for the other bias possibilities. After all, outcomes obtained on the assumption that the bias for heads is  $p$  can be re-interpreted as outcomes for the possibility that the bias is  $1 - p$ , simply by switching the interpretation of success from heads to tails.




Figure 2: Percentages per toss and per bias of agents who assign a high probability to the truth, for specific combinations of  $c$  and  $\epsilon$ ; results are averaged over all simulations.

than EXPL, but EXPL being faster than Bayes’ rule—this picture changed in a social setting, where EXPL appeared to be both more accurate and faster than Bayes’ rule.

It was known that in the kind of individual setting studied in Douven (2013), EXPL users *typically* had a lower Brier score than Bayesians but still tended to have a slightly higher *average* Brier score because *when* they incurred a higher Brier score than Bayesians, it tended to be *much* higher. That is because EXPL users tend to pick up the lead in the evidence faster than Bayesians do, which however has as a downside that they are more prone to follow *misleading* evidence. As Douven and Wenmackers (2017) showed, in a social setting the risk of being pushed off-track by misleading evidence is much diminished due to the averaging mechanism that is called at each step. Specifically, if an agent receives misleading evidence, then the averaging part still brings about a kind of regression to the mean, thereby preventing the agent from straying too far from the true bias probability. (For a more detailed explanation, see Douven & Wenmackers, 2017, Sect. 5.)

### 3 Further generalizing group learning

Our concern, rather than with a direct comparison between Bayes’ rule and EXPL, is with how to epistemically organize a research group in order for that group to strike the best balance between being fast and being accurate. Part of the recommendation to a group might be to adopt Bayes’ rule, or EXPL, as their members’ way of accommodating direct evidence. But, first, we saw in Section 2.1 that there exist other formalizations of explanatory reasoning, and these should certainly also be considered in a comparison. Second, Douven and Wenmackers (2017) let the agents in their simulations average via linear pooling, but one can think of other ways of merging probability functions, some of which might offer advantages for the purpose of striking the best balance between speed and accuracy. Finally, the same authors measure distances between agents’ probability functions by adding up proposition-wise absolute differences between probabilities; this is the so-called Manhattan distance. Here, too, there are variants which could be tried and which, if adopted by a group of agents to determine epistemic peerhood, might help that group to do better, epistemically speaking.

While the linear pooling of probability functions assumed by Douven and Wenmackers (2017) has been advocated by various authors (e.g., Aczél & Wagner, 1980; Lehrer & Wagner, 1981), others have expressed a preference for *geometric pooling* (e.g., Genest, 1984), in which ones takes the geometric instead of the arithmetic average of probabilities (and then renormalizes, which becomes necessary). Formally, the geometric average of numbers  $x_1, x_2, \dots, x_n$  is  $\sqrt[n]{x_1 x_2 \cdots x_n}$ . So, instead of setting, for any update  $u$ , and with  $\Pr_k^{u_{\text{evid}}}(\cdot) \in X(\Pr_j^{u_{\text{evid}}}(\cdot))$  and  $\Psi = \{\psi_i\}_{1 \leq i \leq n}$  a set of jointly exhaustive and

mutually exclusive self-consistent propositions,

$$\Pr_j^{\mathcal{U}_{av}}(\psi_i) = \frac{\sum_k \Pr_k^{\mathcal{U}_{evid}}(\psi_i)}{|X(\Pr_j^{\mathcal{U}_{evid}}(\cdot))|}, \quad \text{for all } i: 1 \leq i \leq n,$$

the proposal is to set

$$\Pr_j^{\mathcal{U}_{av}}(\psi_i) = a \sqrt[m]{\prod_k \Pr_k^{\mathcal{U}_{evid}}(\psi_i)}, \quad \text{for all } i: 1 \leq i \leq n,$$

with  $m = |X(\Pr_j^{\mathcal{U}_{evid}}(\cdot))|$ , and with  $a = 1 / \left( \sum_{i=1}^n \sqrt[m]{\prod_k \Pr_k^{\mathcal{U}_{evid}}(\psi_i)} \right)$ , to ensure that the probabilities assigned to the members of  $\Psi$  sum to 1 again.

For numbers  $x_1, x_2, \dots, x_n$ , their harmonic average equals

$$\left( \frac{x_1^{-1} + x_2^{-1} + \dots + x_n^{-1}}{n} \right)^{-1}.$$

We are not aware of any recommendations in the literature to aggregate probability functions via the harmonic average—“harmonic pooling,” as one might call this aggregation procedure—nor can we think of good a priori reasons for using this procedure. However, neither can we rule out a priori that, for the purpose of making group learning fast and accurate, using the harmonic average has an edge over using either the arithmetic or the geometric average. Thus, below we will also consider the following definition of an average probability function:

$$\Pr_j^{\mathcal{U}_{av}}(\psi_i) = a \left( \frac{\sum_k \Pr_k^{\mathcal{U}_{evid}}(\psi_i)^{-1}}{|X(\Pr_j^{\mathcal{U}_{evid}}(\cdot))|} \right)^{-1}, \quad \text{for all } i: 1 \leq i \leq n,$$

with  $a$  again ensuring that the probabilities sum to 1.

As for measuring distances between probability functions—in order to determine who is an epistemic peer and who is not—the Manhattan distance used by Douven and Wenmackers (2017) is but one possibility. There are many alternatives. Here, we focus on just two, but readers are invited to use the computer code in the Supplementary Materials to experiment with other distances (see below for details). The two to be considered are the Euclidean metric, according to which

$$\Delta(\Pr_i(\cdot), \Pr_j(\cdot)) := \sum_{k=1}^n (\Pr_i(\psi_k) - \Pr_j(\psi_k))^2,$$

and the Kullback–Leibler (KL) divergence, according to which

$$\Delta(\Pr_i(\cdot), \Pr_j(\cdot)) := \sum_{k=1}^n \Pr_i(\psi_k) \ln \left( \frac{\Pr_i(\psi_k)}{\Pr_j(\psi_k)} \right).$$

Note that, supposing the KL divergence, because this is not symmetric it will not generally hold that  $\Delta(\Pr_i(\cdot), \Pr_j(\cdot))$  is equal to  $\Delta(\Pr_j(\cdot), \Pr_i(\cdot))$ . However, in the present context, the asymmetry may actually be an advantage, given that, in reality, respecting someone as one’s epistemic peer is not always reciprocated.

We can treat the extension of the HK model from Douven and Wenmackers (2017) as a template, where we can now choose to plug in one of four evidential update rules: Bayes’ rule, EXPL, Good’s rule, and Popper’s rule; one of three averaging rules: linear pooling, geometric pooling, and harmonic pooling; and one of three rules for determining epistemic peerhood: the Manhattan metric, the Euclidean metric, and the KL divergence. Each of the resulting combinations of rules then still only




Figure 3: Average Brier scores for six bias possibilities, for the combination of EXPL, the Manhattan metric, and geometric pooling ( $p = x$  indicates that the chance for each of the coins to come up heads equals  $x$ ).

yields a schema, with slots for  $c$  and  $\epsilon$  still to be filled in. Does it matter which combination of rules we choose, and what values we set for  $c$  and  $\epsilon$ ?

It will probably come as no surprise that the answer is positive. That this is so is easy to verify by rerunning the simulations reported in Douven and Wenmackers (2017), now for other instances of the template. By way of illustration, consider the plots in Figure 3, which graphically summarize results from simulations in which linear pooling was replaced by geometric pooling, but which were otherwise the same as the simulations underlying Figure 1. We see that this seemingly small difference led to a *completely* different picture. Now, Bayesians perform very well on the count of accuracy, supposing they are open to social influence, although, on the condition of a greater openness to such influence (even if only slightly greater), EXPL users who assign *some* (not-too-large) bonus for explanatory goodness do well, too. By way of further illustration, Figure 4 shows that we obtain again a somewhat different picture if harmonic pooling is assumed and we replace the Manhattan metric by the Euclidean one.<sup>4</sup>

But accuracy—measured by the Brier score here—is only one dimension of comparison, speed being the other one, and the goal, ultimately, is to strike the best balance between them. More importantly, we are only considering an illustration here, which by no means provides a systematic answer to our question of how best to organize a group epistemically (which update rule to recommend to the group, which principles to recommend for regulating social influence on belief states).

To forestall misunderstanding, it is to be noted that, in answering this question, we do not have to solve a multi-objective optimization problem, with speed and accuracy as objectives (Deb, 2001; Qian, Yu, & Zhou, 2013; Rakshit & Konar, 2015). As mentioned repeatedly, we are thinking of *striking the best balance between speed and accuracy* as our objective. In doing so, we shall be assuming that what

<sup>4</sup>The color scaling was set separately for the plots in Figures 1, 3, and 4. Using the same scale in the three figures would have rendered the differences in Figures 3 and 4 difficult to discern.




Figure 4: Average Brier scores for six bias possibilities, for the combination of EXPL, the Euclidean metric, and harmonic pooling ( $p = x$  indicates that the chance for each of the coins to come up heads equals  $x$ ).

counts as “best” depends to an *equal* degree on speed and accuracy. This will not always be the case, of course, but for now, it serves as a useful simplifying assumption; we come back to this in the General Discussion. While by thinking of the task before us as a single-objective optimization problem we are avoiding some complications, it remains the case that the search space is mixed, with both discrete (update rule, distance, pooling method) and continuous variables (values for  $c$  and  $\epsilon$ ). This on its own makes analytical methods difficult to apply. A posteriori methods, specifically computational methods, have been applied with greater success to such problems. In the next section, we turn to these methods, taking our cue in particular from the method of evolutionary computing (Holland, 1975; Koza, 1992; Bäck, 1996; Yu, Yao, & Zhou, 2012).

## 4 Evolutionary computing and optimal group learning

The optimization technique to be used in the following is a specific form of evolutionary computing, commonly referred to as “agent-based optimization.” Like evolutionary computing in general, the technique takes its inspiration from principles of biological evolution. In this technique, different types of agents represent different solutions to a given problem. These solutions are scored with respect to whichever the objective or objectives are, from which in turn the agents’ “fitness” levels are derived. In the next step, the fittest agents are selected, either deterministically or stochastically, in order to be retained or allowed to reproduce (or both) in the next round, where the competition for survival and reproduction begins anew. This can be repeated several times over, until some pre-determined criterion is met (see, e.g., Barbati, Bruno, & Genovese, 2012; also various papers in Czarnowski, Jędrzejowicz, & Kacprzyk, 2013).

Rather than individual agents, we will be considering *groups* of agents, where these groups differ by having adopted different epistemic principles of the kind looked at in the previous sections. One could consider groups whose members use (partly) different epistemic principles, in order to investigate whether such epistemic in-group diversity might carry evolutionary advantages, but we set this possibility aside for future work. In the following, all agents in a group assume the same set of epistemic principles.

Sometimes in evolutionary computing, the pre-determined stopping criterion is a minimally acceptable fitness level to be attained by some agent (or more generally, by some solution to the target problem). For our present concerns, there does not appear to be such a minimal fitness level that could be set a priori. We would *like* groups to get to the truth immediately and unexceptionally, but that is an unrealistic goal to aim at. What *is* a realistic goal is not clear from the outset but rather something we hope to learn from the simulations. So we decided to use the most straightforward stopping criterion to be found in the evolutionary computing literature, to wit, stopping after a fixed number of generations. However, the number we chose—fifty—was established on the basis of some pilot runs that suggested that this number of generations would suffice to achieve satisfactory convergence, in that solutions in the last generation would deviate only slightly from each other. (This was confirmed by the actual runs we carried out, as will be seen below.)

Although all evolutionary algorithms build on the same ideas of variation and selection, they can differ greatly in their details, most notably concerning the extent to which they rely on the basic evolutionary operations of selection, crossover, and mutation as well as the extent to which such operations are applied stochastically or deterministically. The algorithm we used involved both selection and crossover, and to a lesser extent mutation, where the selection proceeded deterministically but the crossover and mutation stochastically. Selection was done on the basis of speed and accuracy, the latter operationalized in terms of the average Brier score incurred by a group, and the former in terms of the first time a majority of the group assigned a high probability (understood as a probability above .9) to the true hypothesis. Both criteria factored equally in determining fitness values.

The simulations we ran were similar to the ones reported in Douven and Wenmackers in that the agents in any one group received binomial data from a common source, where however the data one agent received were independent from those received by any other agent; we can think of the situation as one in which each agent repeatedly flips her own coin, even though all coins have the same (initially unknown) bias, the shared task of the group being to find out what that bias is. For concreteness, we describe the data in the following as being the outcomes of coin flips indeed.

More exactly, then, the algorithm that was used consists of these five steps:

- Step 1. Create a starting population of groups each consisting of 50 agents, such that all combinations of update rules (Bayes' rule, EXPL, Good's rule, Popper's rule), distances (Euclidean, Manhattan, KL divergence), and pooling methods (linear pooling, geometric pooling, harmonic pooling) are represented by exactly one group, yielding a total of 36 groups; for the groups consisting of non-Bayesians updaters (users of EXPL, Good's rule, or Popper's rule), choose randomly per group an explanatory bonus value from a standard uniform distribution; and for each group, choose randomly a value for  $\epsilon$  also from the standard uniform distribution.
- Step 2. For every agent in every group, generate a sequence of 500 coin flips, where all coins have the same randomly chosen bias (the choices being limited in the exact way as in Douven & Wenmackers, 2017, meaning that all possible biases are multiples of 0.1). Let agents update their probabilities after each toss according to whichever combination of rules their group has



adopted—using the generalized HK model described earlier<sup>5</sup>—and after each update, calculate Brier scores for all agents. Also, register after each update the percentage of agents who assign a probability greater than .9 to whichever the true bias hypothesis is in this run. For each group, repeat this scoring process 25 times; so choose anew a bias for the coins, let the agents update on 500 coin flips, and score each agent with respect to accuracy (Brier score) and speed (probability greater than .9 assigned to the truth), and keep doing this until all agents have gone through 25 such rounds of scoring. Then calculate the group’s average Brier score (i.e., the total Brier score incurred by the group’s members, averaged over the 25 runs) and also the average toss number after which a majority of agents in the group assign a probability greater than .9 to the truth (the average again being over the 25 runs).

Step 3. Rank the 36 groups in the population twice, once according to their average Brier score and once according to their average speed, with lower ranks being better in both cases. Then add up the ranks and select the 18 groups with the lowest resulting sums (following, e.g., Baker, 1985). Draw (with replacement) from these “fittest” groups 18 pairs and then select randomly and independently from each pair an update rule, a distance, and a pooling method; draw a value for the explanation bonus  $c$  from a normal distribution truncated between 0 and 1 with as mean the mean of the  $c$  values of the members of the pair and as standard deviation the standard deviation of the  $c$  values of the whole previous generation; and determine the value for  $\varepsilon$  in the exact same way. Put the results of these selections together to obtain 18 new “offspring” groups.

Step 4. Form the next generation of groups by conjoining the 18 best groups from the parent generation and their offspring (another 18 groups) formed as explained in Step 3.

Step 5. Iterate Steps 2–4.

Figure 5 gives a schematic representation of the above steps.

We have two comments on this algorithm. First, as mentioned above, selection occurs deterministically while variation occurs stochastically. Given the way the values of  $c$  and  $\varepsilon$  are picked for the offspring, an element of random mutation is built into the variation part. In the piloting stage of the present research, we experimented with stochastic selection as well, in particular with an algorithm in which a group’s rank on the basis of the scoring determined its chances of being selected for producing offspring. Because this did not lead to results qualitatively different from the one’s obtained with deterministic selection but did slow down the computations by quite a bit, we went for deterministic selection in the full simulations.

Second, the simulations were conducted using the new computer language Julia (Bezanson et al., 2017). Julia code reads almost like pseudo-code. Therefore, readers interested in a more detailed description of the simulations are invited to consult the Jupyter notebook containing the code, available at [this link](#). The code will also allow readers to try out variant update rules, distances, and aggregation functions.<sup>6</sup> For example, the distances we used are from the Distances.jl package for Julia, which currently supports a total of 27 distances, a number of which are certainly plausible options for determining peerhood (e.g., the Jensen–Shannon divergence and the Hellinger distance).

<sup>5</sup>To forestall misunderstanding, it is to be noted that the agents in the model all belong to the same group, so that the algorithm only enables intra-group communication, and not inter-group communication.

<sup>6</sup>There are at least two ways to run the notebook. First, one can upload the notebook to JuliaBox, at <https://juliabox.com>, which is a free service. Second, one can run the notebook locally, on one’s own computer, which of course requires Julia to be installed and preferably also JupyterLab, both of which are free and open source software. Installation instructions




Figure 5: Schematic representation of how the evolutionary algorithm creates a new generation of groups. (Solid arrows are meant to suggest that groups move from one stage to the next, and the dotted arrow that, while it is not groups themselves that move to a next stage, their genetic material does.)

As said, one run consisted of 50 generations where, except for the starting generation each generation was the result of the process described in Steps 2–4. We planned to start by conducting 15 runs, and to continue with additional runs depending on how much variance we were to find in the first 15.

Turning now to an evaluation of the results we obtained, we get a first impression of what they looked like from Figures 6 and 7. These figures show the outcomes from a single, randomly chosen run. Figure 6 consists of three series of 50 stacked histograms, where each histogram shows the number of tokens of each update rule type (top plot), distance type (middle plot), and pooling type (bottom plot) present in the corresponding generation. It is clear from this figure that after 19 generations, all groups already agree on a best combination of update rule, distance, and pooling method, to wit, EXPL, the Euclidean metric, and harmonic pooling. Also, Figure 7 shows that there is large agreement among the groups about the best values for  $c$  and  $\epsilon$ , which are both in the higher range. From Figure 8 it can be seen that how the population of groups evolves has a noticeable impact on how accurate they are, on average, as well as on how fast they are, on average.

While the aforementioned figures show the results of a randomly selected run, we obtained very similar results in the other runs. Notably, we found a strong convergence not only in that all runs showed all groups to endorse, after 20 to 30 generations, a combination of EXPL, the Euclidean metric, and the harmonic pooling method, but also in that, in each run, all groups came to hold values for both  $c$  and  $\epsilon$  typically in the vicinity of .9. More specifically still, pooling the results from all runs, we found, averaging over all last-generation groups, a  $c$  value of .94 with a standard deviation of only 0.04, and an  $\epsilon$  value of .91 with a standard deviation of 0.05. These values hardly differ from the average values of the top groups in the last generations from the 15 runs:  $.95 (\pm 0.03)$  for  $c$ , and  $.89 (\pm 0.07)$  for  $\epsilon$ . In view of these results, we refrained from conducting runs in addition to the 15 we started with.

To see the effect the evolutionary process had on the criteria of interest, we note that the average total Brier score incurred by the groups in the first generation was  $6468.19 (\pm 730.07)$  while that incurred by the groups in the last generation was  $616.08 (\pm 114.59)$ . To make these results somewhat comparable

---

for these applications are different for different operating systems and can be found at <https://julialang.org/downloads> and <https://github.com/jupyterlab/jupyterlab>, respectively.

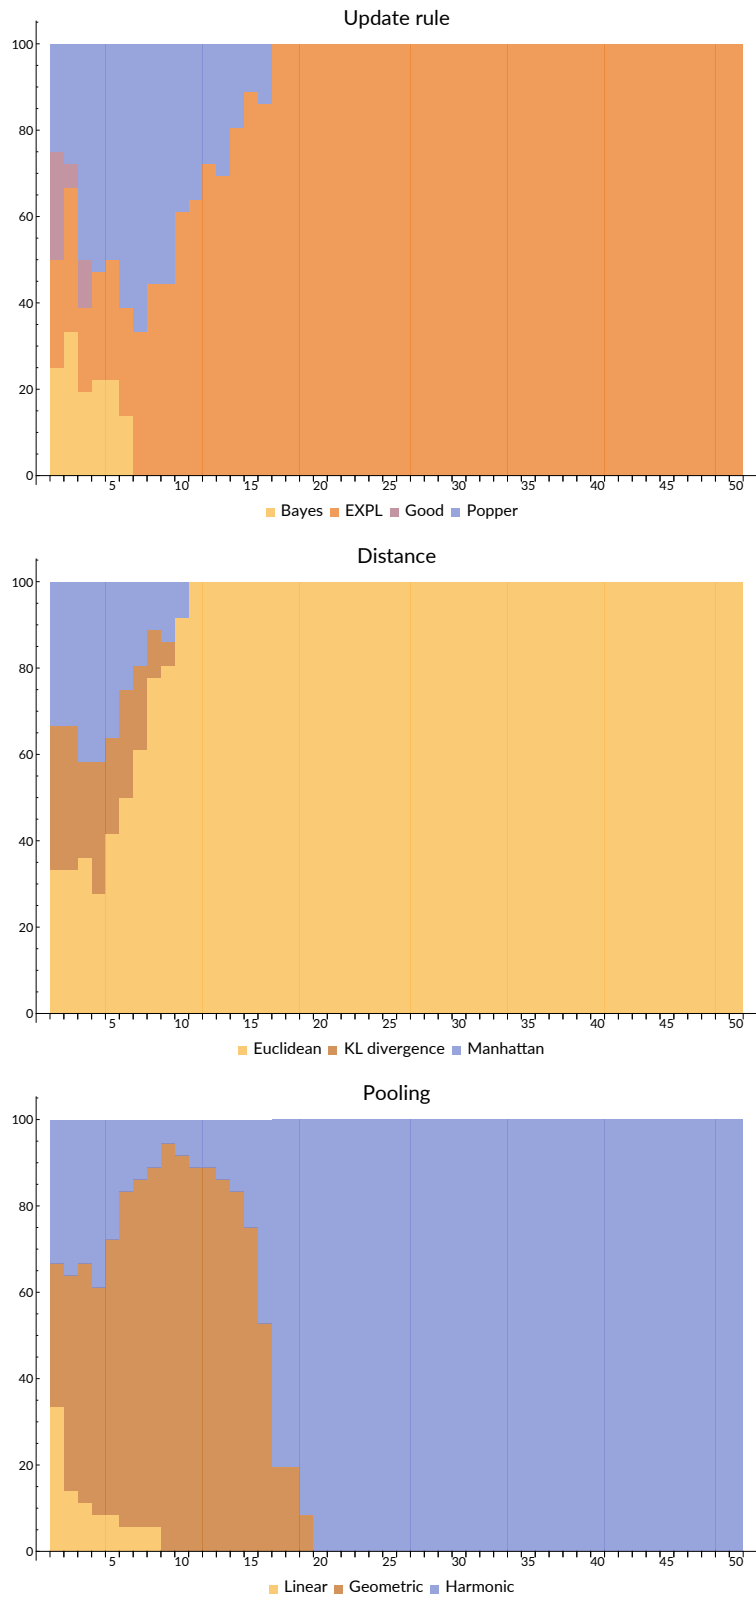


Figure 6: For a randomly chosen run, per-generation counts of rule types (top), distance types (middle), and pooling types (bottom); generation is represented on the horizontal axis, count on the vertical axis. See the text for further explanation.




Figure 7: Average bonus and  $\epsilon$  values for the 50 generations in the same run that generated the data shown in Figure 6; generation is represented on the horizontal axes, values are represented on the vertical axes. Error bars indicate one standard deviation from the mean.

with those reported in Section 2.2, note that this amounts to an average Brier score *per agent* of 129.36 for first-generation groups, and 12.32 for last-generation groups. Similarly, the first toss number at which a majority of agents in a group assigned a high probability to the truth was, on average, 153.60 ( $\pm 6.20$ ) for groups in the first generation, and 31.47 ( $\pm 1.16$ ) for groups in the last generation.

We can bring the significance of the overall outcome into further relief by running the procedure for determining fitness levels used in the evolutionary computations on a group of non-interacting Bayesians (agents updating via Bayes' rule but not letting themselves be influenced by the opinions of their peers) and, for comparison, on a group of agents that has adopted the setting that came out best in the foregoing, so a combination of EXPL as an update rule, with  $c = 0.95$ ; the Euclidean metric for determining peerhood, with  $\epsilon = .89$  as determining the criterion for peerhood; and the harmonic pooling function for aggregating probability functions. For either group, we can (i) draw a random bias hypothesis; (ii) generate for each agent her own sequence of 500 tosses with the drawn bias; (iii) let each agent consecutively update on those tosses and, for those in the second group, also take into account her peers' opinions in the way prescribed by the said epistemic principles; (iv) register the Brier scores for each agent after each update; (v) register the first toss after which a majority of the group assigns a probability above .9 to whichever bias hypothesis was drawn in step (i); and (vi) repeat the foregoing




Figure 8: Average Brier scores and average first toss after which the majority assigns a high probability to the truth, for the 50 generations in the same run that generated the data shown in Figures 6 and 7; generation is represented on the horizontal axes, average score and speed are represented on the vertical axes. Error bars indicate one standard deviation from the mean.

five steps 25 times, just like in the evolutionary computations. We ran these simulations 10 times, in which the average total Brier score was 3798.82 ( $\pm 347.43$ ) for the Bayesians and 524.91 ( $\pm 59.57$ ) for the group using the “best” settings, and in which the average toss at which a majority assigned a probability above .9 to the truth was 118.58 ( $\pm 10.00$ ) for the Bayesians and 29.90 ( $\pm 4.12$ ) for the other group. In other words, by using the settings the evolutionary computations showed to be optimal, a group can on average achieve a more than sevenfold increase in accuracy and at the same time an almost fourfold increase in speed, relative to a group of strict Bayesian updaters.

## 5 General discussion

Our overarching research question was whether we can recommend best epistemic practices to groups of agents who independently receive evidence from a common source. To address this question, we used the method of agent-based optimization. Assuming that a group is interested in striking the best balance between speed and accuracy, and making certain assumptions about the evidence-generating process, we found that we can make a clear recommendation in the form of a unique combination of update rule, distance measure, and pooling method, together with an indication of the optimal range for the  $c$  and  $\varepsilon$  values, the former indicating the explanatory bonus, the latter willingness to count others as peers.

This result was reached by assuming that speed and accuracy mattered equally. As said, however, this assumption is not guaranteed to hold. In some situations, we may be willing to sacrifice quite a bit of accuracy for the benefit of greater speed, while in other situations, the opposite may hold. The code that was used for the simulations, and that was linked to earlier, allows interested readers to verify that giving different weights to speed and accuracy can lead to vastly different outcomes. For instance, rerunning the simulations with all the weight given to accuracy consistently led to a combination of Popper’s rule, the Euclidean metric, and geometric pooling as the optimal combination, with values for  $c$  and  $\varepsilon$  in the mid-to-high range. Apart from the relative importance of speed and accuracy, the evidence in our simulations also had properties which do not hold generally—for instance, if the data really concern coin tosses, one might be more likely to encounter biases close to .5 than ones more to either extreme—and varying the evidence-generating process can affect the outcomes as well.<sup>7</sup>

Note, however, that nothing follows that was not already recognized by cognitive psychologists working on theories of ecological rationality. According to these researchers, the idea—prominent among philosophers—that rationality can be captured by a number of universally valid principles has been an illusion all along. As argued by Gigerenzer and his various collaborators (e.g., Gigerenzer & Goldstein, 1996; Gigerenzer et al., 1999; Gigerenzer, 2000; Goldstein & Gigerenzer, 2002; Gigerenzer, Hertwig, & Pachur, 2011; Todd & Gigerenzer, 2012; Todd & Brighton, 2016; Schurz & Hertwig, 2018), and also by Elqayam (2011, 2012), we can understand the notion of rationality only if we take into account (i) the various ways—biological and cognitive—in which humans are limited, and (ii) the environment or environments in which humans are to operate and to interact. Biological, cognitive, and environmental constraints can all vary greatly from one individual to another and even from one

---

<sup>7</sup> An anonymous referee conjectured that number of hypotheses mattered, too. Rerunning the simulations, now for five hypotheses, gave indeed different results than those reported in the main text. In particular, while the winning combination of update rule, metric, and pooling method (EXPL/Euclidean/harmonic) was still the winner, in 40 percent of the last generations we now found the combination of Good’s rule, the Euclidean metric, and geometric pooling. The mean values for  $c$  and  $\varepsilon$  were .81 and .95, respectively. It may be debatable whether five, or even eleven, hypotheses is a realistic number of hypotheses to consider in a scientific setting (Douven & Uffink, 2012; Douven & Rott, 2018), and running simulations for, say, 100 hypotheses might yield results much more different still from those we obtained. Needless to say, however, such simulations would be computationally expensive.

moment to another. This led Gigerenzer, Elqayam, and others to conclude that whether someone reasons or acts rationally depends on whether the reasoning or behavior helps the person to achieve whatever goal or goals she pursues in whichever environment she is in, given the cognitive resources available to her in that environment.

The way we proceeded in our simulations is in line with Arkes, Gigerenzer, and Hertwig's (2016) guidelines for investigating the ecological rationality of a cognitive strategy (e.g., a group-learning strategy; see also Anderson, 1990, 1991). According to these authors, such an investigation should consist of two steps. In the first step, one should determine (i) the goal of an individual or a group; (ii) the strategies available to the individual or group for achieving that goal; and (iii) the structural properties of the individual's or group's environment. And then, in the second step, one should determine which strategy or strategies are most likely to help the individual or group achieve the goal in the given environment. In our case, the goal was, by stipulation, the group's wish to find an optimal trade-off between speed and accuracy; the available strategies were the various combinations of update rule, distance, and pooling method; and the relevant properties of the environment were defined by the data-generating process. We relied on an evolutionary algorithm to accomplish the second step.

In short, the answer to our research question is that we *can* make recommendations for how to organize research groups epistemically, but these recommendations will depend on what type of phenomenon or process the group is studying and on what the research aim of a group is. In view of the literature on ecological rationality, a stronger conclusion could not have been expected.<sup>8</sup>

## References

- Aczél, J. & Wagner, C. (1980). A characterization of weighted arithmetic means. *SIAM Journal on Algebraic and Discrete Methods*, 1, 259–260.
- Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50, 510–530.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale NJ: Lawrence Erlbaum Associates.
- Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, 14, 471–517.
- Arkes, H. R., Gigerenzer, G., & Hertwig, R. (2016). How bad is incoherence? *Decision*, 3, 20–39.
- Bäck, T. (1996). *Evolutionary algorithms in theory and practice*. Oxford: Oxford University Press.
- Baker, J. E. (1985). Adaptive selection methods for genetic algorithms. In J. J. Grefenstette (ed.), *Proceedings of the first international conference on genetic algorithms and their applications* (pp. 100–111). Hillsdale NJ: Lawrence Erlbaum Associates.
- Baral, C. (2000). Abductive reasoning through filtering. *Artificial Intelligence*, 120, 1–28.
- Barbati, M., Bruno, G., & Genovese, A. (2012). Applications of agent-based models for optimization problems: A literature review. *Expert Systems with Applications*, 39, 6020–6028.
- Bes, B., Sloman, S., Lucas, C. G., & Raufaste, E. (2012). Non-Bayesian inference: Causal structure trumps correlation. *Cognitive Science*, 36, 1178–1203.
- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, 59, 65–98.
- Boutilier, C. & Becher, V. (1995). Abduction as belief revision. *Artificial Intelligence*, 77, 43–94.
- Boyd, R. (1984). The current status of scientific realism. In J. Leplin (ed.), *Scientific realism* (pp. 41–82). Berkeley CA: University of California Press.

---

<sup>8</sup>I am greatly indebted to three anonymous referees for very constructive comments on a previous version of this paper. Thanks also to Christopher von Bülow for valuable comments.

- Boyd, R. (1990). Realism, approximate truth, and philosophical method. In C. Wade Savage (ed.), *Scientific theories* (pp. 355–391). Minneapolis MN: University of Minnesota Press.
- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78, 1–3.
- Costello, F. & Watts, P. (2016). People’s conditional probability judgments follow probability theory (plus noise). *Cognitive Psychology*, 89, 106–133.
- Costello, F. & Watts, P. (2018). Invariants in probabilistic reasoning. *Cognitive Psychology*, 100, 1–16.
- Crosscombe, M. & Lawry, J. (2016). A model of multi-agent consensus for vague and uncertain beliefs. *Adaptive Behavior*, 24, 249–260.
- Czarnowski, I., Jędrzejowicz, P., & Kacprzyk, J. (2013). *Agent-based optimization*. Berlin: Springer.
- Darwiche, A. (2009). *Modeling and reasoning with Bayesian networks*. Cambridge: Cambridge University Press.
- Deb, K. (2001). *Multi-objective optimization using evolutionary algorithms*. Chichester UK: Wiley.
- Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3, 87–98.
- De Langhe, R. (2013). Peer disagreement under multiple epistemic constraints. *Synthese*, 190, 2547–2556.
- De Langhe, R. & Greiff, M. (2010). Standards and the distribution of cognitive labour: A model of the dynamics of scientific activity. *Logic Journal of the IGPL*, 18, 278–294.
- Dittmer, J. C. (2001). Consensus formation under bounded confidence. *Nonlinear Analysis*, 7, 4615–4621.
- Douven, I. (1999). Inference to the best explanation made coherent. *Philosophy of Science*, 66, S424–S435.
- Douven, I. (2010). Simulating peer disagreements. *Studies in History and Philosophy of Science*, 41, 148–157.
- Douven, I. (2013). Inference to the best explanation, Dutch books, and inaccuracy minimisation. *Philosophical Quarterly*, 69, 428–444.
- Douven, I. (2017). Inference to the best explanation: What is it? And why should we care? In T. Poston & K. McCain (eds.), *Best explanations: New essays on inference to the best explanation* (pp. 4–22). Oxford: Oxford University Press.
- Douven, I. (2019). Can the evidence for explanatory reasoning be explained away? *IEEE Transactions on Cognitive and Developmental Systems*, in press.
- Douven, I. & Kelp, C. (2011). Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis*, 75, 271–283.
- Douven, I. & Mirabile, P. (2018). Best, second-best, and good-enough explanations: How they matter to reasoning. *Journal of Experimental Psychology: Language, Memory, and Cognition*, 44, 1792–1813.
- Douven, I. & Riegler, A. (2010). Extending the Heggelmann–Krause model I. *Logic Journal of the IGPL*, 18, 323–335.
- Douven, I. & Rott, H. (2018). From probabilities to categorical beliefs: Going beyond toy models. *Journal of Logic and Computation*, 28, 1099–1124.
- Douven, I. & Schupbach, J. N. (2015a). The role of explanatory considerations in updating. *Cognition*, 142, 299–311.
- Douven, I. & Schupbach, J. N. (2015b). Probabilistic alternatives to Bayesianism: The case of explanationism. *Frontiers in Psychology*, 6, doi: 10.3389/fpsyg.2015.00459.
- Douven, I. & Uffink, J. (2012). Quantum probabilities and the conjunction principle. *Synthese*, 184, 109–114.



- Douven, I. & Wenmackers, S. (2017). Inference to the best explanation versus Bayes' rule in a social setting. *British Journal for the Philosophy of Science*, 68, 535–570.
- Earman, J. (1992). *Bayes or bust?* Cambridge MA: MIT Press.
- Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (ed.), *Formal representation of human judgment* (pp. 17–52). New York: Wiley.
- Elqayam, S. (2011). Grounded rationality: A relativist framework for normative rationality. In K. I. Manktelow, D. E. Over, & S. Elqayam (eds.), *The science of reason* (pp. 397–420). Hove UK: Psychology Press.
- Elqayam, S. (2012). Grounded rationality: Descriptivism in epistemic context. *Synthese*, 189, 39–49.
- Fischhoff, B. & Lichtenstein, S. (1978). Don't attribute this to Reverend Bayes. *Psychological Bulletin*, 85, 239–243.
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., & Rubin, D. (2013). *Bayesian data analysis*. New York: Chapman and Hall/CRC.
- Genest, C. (1984). A characterization theorem for externally Bayesian groups. *Annals of Statistics*, 12, 1100–1105.
- Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. New York: Oxford University Press.
- Gigerenzer, G. & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650–669.
- Gigerenzer, G., Hertwig, R., & Pachur, T. (2011). *Heuristics: The foundations of adaptive behavior*. New York: Oxford University Press.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group (1999). *Simple heuristics that make us smart*. Oxford: Oxford University Press.
- Glass, D. H. (2007). Coherence measures and inference to the best explanation. *Synthese*, 157, 275–296.
- Glass, D. H. (2012). Inference to the best explanation: Does it track truth? *Synthese*, 185, 411–427.
- Glass, D. H. (2018a). An evaluation of probabilistic approaches to inference to the best explanation. *International Journal of Approximate Reasoning*, 103, 184–194.
- Glass, D. H. (2018b). Coherence, explanation, and hypothesis selection. *British Journal for the Philosophy of Science*, in press.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Goldstein, D. G. & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109, 75–90.
- Good, I. J. (1960). Weight of evidence, corroboration, explanatory power, information and the utility of experiment. *Journal of the Royal Statistical Society B*, 22, 319–331.
- Gottifredi, S., Tamargo, L. H., García, A. J., & Simari, G. R. (2018). Arguing about informant credibility in open multi-agent systems. *Artificial Intelligence*, 259, 91–109.
- Gribbin, J. (2002). *Science: A history*. London: Penguin.
- Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science*, 7, 217–229.
- Hahn, U., Hansen, J. U., & Olsson, E. J. (2019). Truth tracking performance of social networks: How connectivity and clustering can make groups less competent. *Synthese*, in press.
- Hegselmann, R. & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis, and simulations. *Journal of Artificial Societies and Social Simulation*, 5, <http://jasss.soc.surrey.ac.uk/5/3/2.html>.
- Hegselmann, R. & Krause, U. (2005). Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25, 381–405.



- Hegselmann, R. & Krause, U. (2006). Truth and cognitive division of labor: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation*, 9, <http://jasss.soc.surrey.ac.uk/9/3/10.html>.
- Hegselmann, R. & Krause, U. (2009). Deliberative exchange, truth, and cognitive division of labour: A low-resolution modeling approach. *Episteme*, 6, 130–144.
- Holland, J. (1975). *Adaptation in natural and artificial systems*. Cambridge MA: MIT Press.
- Howson, C. & Urbach, P. M. (1989). *Scientific reasoning: The Bayesian approach*. La Salle IL: Open Court.
- Jacobmeier, D. (2004). Multidimensional consensus model on a Barabási–Albert network. *International Journal of Modern Physics C*, 16, 633–646.
- Jeffrey, R. (1965). *The logic of decision*. New York: McGraw-Hill.
- Johnston, A. M., Johnson, S. G. B., Koven, M. L., & Keil, F. C. (2017). Little Bayesians or little Einsteins? Probability and explanatory virtue in children’s inferences. *Developmental Science*, 20, e12483, <https://doi.org/10.1111/desc.12483>.
- Kern-Isberner, G. & Lukasiewicz, T. (2004). Combining probabilistic logic programming with the power of maximum entropy. *Artificial Intelligence*, 157, 139–202.
- Kitcher, P. (1993). *The advancement of science*. Oxford: Oxford University Press.
- Konolige, K. (1992). Abduction versus closure in causal theories. *Artificial Intelligence*, 53, 255–272.
- Korb, K. B. & Nicholson, A. N. (2004). *Bayesian artificial intelligence*. London: Chapman and Hall.
- Koslowski, B. (2018). Abductive reasoning and explanation. In L. J. Ball & V. A. Thompson (eds.), *International handbook of thinking and reasoning* (pp. 366–382). London: Routledge.
- Koslowski, B., Marasia, J., Chelenza, M., & Dublin, R. (2008). Information becomes evidence when an explanation can incorporate it into a causal framework. *Causal Development*, 23, 472–487.
- Koza, J. (1992). *Genetic programming: On the programming of computers by means of natural selection*. Cambridge MA: MIT Press.
- Kruschke, J. K. (2015). *Doing Bayesian data analysis* (2nd ed.). Amsterdam: Academic Press.
- Kuipers, T. A. F. (2000). *From instrumentalism to constructive realism*. Dordrecht: Kluwer.
- Kuipers, T. A. F. (2001). *Structures in science*. Dordrecht: Kluwer.
- Kuipers, T. A. F. (2014). Empirical progress and nomic truth approximation revisited. *Studies in History and Philosophy of Science*, 46, 64–72.
- Kuipers, T. A. F. (2019). *Nomic truth approximation revisited*. Basel: Springer.
- Kummerfeld, E. & Zollman, K. J. S. (2016). Conservatism and the scientific state of nature. *British Journal for the Philosophy of Science*, 82, 956–968.
- Legare, C. H. & Lombrozo, T. (2014). Selective effects of explanation on learning during early childhood. *Journal of Experimental Child Psychology*, 126, 198–212.
- Lehrer, K. & Wagner, C. (1981). *Rational consensus in science and society*. Dordrecht: Reidel.
- Lewis, D. K. (1980). A subjectivist’s guide to objective chance. In R. C. Jeffrey (ed.), *Studies in inductive logic and probability, Vol. II* (pp. 263–293). Berkeley CA: University of California Press.
- Lin, F. & You, J.-H. (2002). Abduction in logic programming: A new definition and an abductive procedure based on rewriting. *Artificial Intelligence*, 140, 175–205.
- Lombrozo, T. (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, 20, 748–759.
- Lombrozo, T. & Gwynne, N. Z. (2014). Explanation and inference: Mechanistic and functional explanations guide property generalization. *Frontiers in Human Neuroscience*, 8, <https://doi.org/10.3389/fnhum.2014.00700>.
- Lorenz, J. (2003). *Mehrdimensionale Meinungsdynamik bei wechselndem Vertrauen*. Diploma thesis, University of Bremen, available at <http://nbn-resolving.de/urn:nbn:de:gbv:46-dipl000000564>.

- Lorenz, J. (2008). Fostering consensus in multidimensional continuous opinion dynamics under bounded confidence. In D. Helbing (ed.), *Managing complexity: Insights, concepts, applications* (pp. 321–334). Berlin/Heidelberg: Springer.
- Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences of the United States of America*, 108, 9020–9025.
- Marks, D. F., & Clarkson, J. K. (1972). An explanation of conservatism in the bookbag-and-pokerchips situation. *Acta Psychologica*, 36, 145–160.
- Mason, W. A., Conrey, F. R., & Smith, E. R. (2007). Situating social influence processes: Dynamic, multidirectional flows of influence within social networks. *Personality and Social Psychology Review*, 11, 279–300.
- Mason, W. A. & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences of the United States of America*, 109, 764–769.
- McMullin, E. (1992). *The inference that makes science*. Milwaukee WI: Marquette University Press.
- Miller, D. T. & McFarland, C. (1991). When social comparison goes awry: The case of pluralistic ignorance. In J. Suls & T. Wills (eds.), *Social comparison: Contemporary theory and research* (pp. 287–313). Hillsdale NJ: Lawrence Erlbaum Associates.
- Nunes, D. & Antunes, L. (2015). Modelling structured societies: A multi-relational approach to context permeability. *Artificial Intelligence*, 229, 175–199.
- Oaksford, M. & Chater, N. (2007). *Bayesian rationality*. Oxford: Oxford University Press.
- Olsson, E. J. (2008). Knowledge, truth, and bullshit: Reflections on Frankfurt. *Midwest Studies in Philosophy*, 32, 94–110.
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton NJ: Princeton University Press.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Pennington, N. & Hastie, R. (1988). Explanation-based decision making: The effects of memory structure on judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 521–533.
- Pennington, N. & Hastie, R. (1992). Explaining the evidence: Tests of the story-model for juror decision making. *Journal of Personality and Social Psychology*, 62, 189–206.
- Pennington, N. & Hastie, R. (1993). Reasoning in explanation-based decision making. *Cognition*, 49, 123–163.
- Phillips, L. D. & Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology*, 72, 346–354.
- Pluchino, A., Latora, V., & Rapisarda, A. (2006). Compromise and synchronization in opinion dynamics. *European Physical Journal B*, 50, 169–176.
- Popper, K. R. (1959). *The logic of scientific discovery*. London: Hutchinson.
- Psillos, S. (1999). *Scientific realism: How science tracks truth*. London: Routledge.
- Qian, C., Yang, Y., & Zhou, Z.-H. (2013). An analysis on recombination in multi-objective evolutionary optimization. *Artificial Intelligence*, 204, 99–119.
- Rakshit, P. & Konar, A. (2015). Differential evolution for noisy multiobjective optimization. *Artificial Intelligence*, 227, 165–189.
- Riegler, A. & Douven, I. (2009). Extending the Heggelmann–Krause model III: From single beliefs to complex belief states. *Episteme*, 6, 145–163.
- Rosenkrantz, R. D. (1992). The justification of induction. *Philosophy of Science*, 59, 527–539.

- Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In epistemic networks, is less really more? *Philosophy of Science*, 84, 234–252.
- Schum, D. A. & Martin, A. W. (1982). Formal and empirical research on cascaded inference in jurisprudence. *Law and Society Review*, 17, 105–151.
- Schurz, G. & Hertwig, R. (2018). Cognitive success: A consequentialist account of rationality and cognition. *Topics in Cognitive Science*, in press.
- Semeshenko, V., Gordon, M. B., & Nadal, J.-P. (2008). Collective states in social systems with interacting learning agents. *Physica A: Statistical Mechanics and its Applications*, 387, 4903–4916.
- Shoham, Y., Powers, R., & Grenager, T. (2007). If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171, 365–377.
- Stone, M. (1961). The opinion pool. *Annals of Mathematical Statistics*, 32, 1339–1342.
- Surowiecki, J. (2004). *The wisdom of crowds*. New York: Anchor Books.
- Tamargo, L. H., Garcia, A. J., Falappa, M. A., & Simari, G. R. (2014). On the revision of informant credibility orders. *Artificial Intelligence*, 212, 36–58.
- Teijeiro, T. & Félix, P. (2018). On the adoption of abductive reasoning for time series interpretation. *Artificial Intelligence*, 262, 163–188.
- Teller, P. (1973). Conditionalization and observation. *Synthese*, 26, 218–258.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331, 1279–1285.
- Todd, P. M. & Brighton, H. (2016). Building the theory of ecological rationality. *Minds and Machines*, 26, 9–30.
- Todd, P. M. & Gigerenzer, G. (2012). *Ecological rationality: Intelligence in the world*. Oxford: Oxford University Press.
- Trpin, B. & Pellert, M. (2018). Inference to the best explanation in uncertain evidential situations. *British Journal for the Philosophy of Science*, in press.
- van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford: Oxford University Press.
- Weisbuch, G., Deffuant, G., Amblard, F., & Nadal, J.-P. (2002). Meet, discuss and segregate! *Complexity*, 7, 55–63.
- Wenmackers, S., Vanpoucke, D., & Douven, I. (2012). Probability of inconsistencies in theory revision: A multi-agent model for updating logically interconnected beliefs under bounded confidence. *European Physical Journal B*, 85, <https://doi.org/10.1140/epjb/e2011-20617-8>.
- Wenmackers, S., Vanpoucke, D., & Douven, I. (2014). Rationality: A social-epistemology perspective. *Frontiers in Psychology*, 5, 581, <https://doi.org/10.3389/fpsyg.2014.00581>.
- Williams, J. J. & Lombrozo, T. (2013). Explanation and prior knowledge interact to guide learning. *Cognitive Psychology*, 66, 55–84.
- Yu, Y., Yao, X., & Zhou, Z.-H. (2012). On the approximation ability of evolutionary optimization with application to minimum set cover. *Artificial Intelligence*, 180–181, 20–33.
- Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74, 574–587.