



HAL
open science

Integrating individual and social learning: accuracy and evolutionary viability

Igor Douven, Gerhard Schurz

► **To cite this version:**

Igor Douven, Gerhard Schurz. Integrating individual and social learning: accuracy and evolutionary viability. Computational and Mathematical Organization Theory, In press, 10.1007/s10588-022-09372-1 . hal-03921574

HAL Id: hal-03921574

<https://cnrs.hal.science/hal-03921574v1>

Submitted on 4 Jan 2023

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/366187191>

Integrating Individual and Social Learning: Accuracy and Evolutionary Viability

Article in *Computational and Mathematical Organization Theory* · December 2022

DOI: 10.1007/s10588-022-09372-1

CITATIONS

0

READS

41

2 authors, including:



Igor Douven

French National Centre for Scientific Research

212 PUBLICATIONS 3,364 CITATIONS

SEE PROFILE

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



Realism [View project](#)



Optimal categorization [View project](#)

Integrating Individual and Social Learning: Accuracy and Evolutionary Viability*

Igor Douven

IHPST / CNRS / Panthéon–Sorbonne University

igor.douven@univ-paris1.fr

Gerhard Schurz

Heinrich Heine University of Düsseldorf

gerhard.schurz@uni-duesseldorf.de

Abstract

Much of what we know, we know thanks to our interactions with others. There is a variety of ways in which we learn from others. We sometimes simply adopt the viewpoints of those we regard as experts, but we also sometimes change our viewpoints in more subtle ways based on the viewpoints of people we regard as our peers. Both forms of social learning have been receiving increasing attention. However, studies investigating how best to combine them, and how to combine the two with individual forms of learning, are still few and far between. This paper looks at ways to integrate various forms of social learning with learning at an individual level within a broadly Bayesian framework. Using agent-based models, we compare the different ways in terms of accuracy of belief states as well as in terms of evolutionary viability. The outcomes of our simulations suggest that agents are best off spending most of their time engaging in social learning, reserving only a limited amount of time for individual learning.

Keywords: accuracy; agent-based modeling; evolutionary computation; bounded confidence model; induction; meta-induction; probability; simulations; social learning.

1 Introduction

Although we are equipped to explore the world on our own, it is no more than a commonplace to acknowledge that much of what we know, we know thanks to others (Boyd & Richerson, 1985; Goldman, 1999; Tomasello, 1999, 2019). We learn from our parents, our teachers, our friends, our colleagues, even from our students and adversaries. There is no one way in which we learn from others. When we consider someone as an expert on a given matter, we may simply align our own opinion on that matter with hers (Gaifman, 1986; van Fraassen, 1989; Goldman, 2001). When, by contrast, we consider someone our “epistemic peer”—someone we deem to be intellectually our equal and to be about equally well informed as we are—her opinion may impact in more subtle ways on ours (Goldman, 2010).

*The code for the simulations and analyses reported in this paper was written in the scientific computing language Julia (Bezanson et al., 2017). Interested readers can download the code from this repository: <https://github.com/IgorDouven/Integrating-Individual-and-Social-Learning>.

Recent years have seen an increase of interest in the study of social learning, leading to both new theoretical work (e.g., Cesa-Bianchi & Lugosi, 2006; Schurz, 2019) and new experimental work (e.g., Darr, Argote & Epple, 1995; Kane, Argote, & Levine, 2005; Harris & Corriveau, 2011; Lorenz et al., 2011; Wood, Kendal, & Flynn, 2012, 2013; Muthukrishna, Morgan, & Henrich, 2016) on social learning strategies. Especially the theoretical work also gave rise to new applications of agent-based modeling (Fortunato, 2004; Pluchino, Latora, & Rapisarda, 2006; Zollman, 2007, 2010, 2011; Hegselmann, 2014, 2020; Hegselmann et al., 2015; Crosscombe & Lawry, 2016; Kummerfeld & Zollman, 2016; Rosenstock, Bruner, & O'Connor, 2017; Douven, 2019, 2022a; O'Connor & Weatherall, 2019; Douven & Hegselmann, 2021, 2022; Schawe, Fontaine, & Hernández, 2021). Although this work has brought a wealth of new insights into the role and importance of social learning, most of the agent-based models that came out of it are focused either on success-independent learning from peers or on success-dependent learning from experts, where this is or is not combined with individual learning. There have been no attempts so far to integrate, in a realistic way, individual learning, learning from peers, and learning from experts. To arrive at a fuller understanding of the significance of social learning, we must look at how the various ways of learning can play together.

More specifically still, we will want to know how to most successfully integrate individual and social learning. Are there any strategies for combining the two forms of learning that will help us arrive more quickly or more reliably at our epistemic goal or goals (typically, related to finding the truth about particular areas of interest)? This question can be naturally thought of from a social engineering perspective. Suppose you are the head of a scientific institute or of the research and development department of a company. Then it will be of considerable interest to know how best to organize the members of your team to optimize its performance and, for instance, maximize its chances of discovering the cause of a given disease, or developing a new vaccine, or more generally, achieving whatever the overarching research goal is. Equipped with such knowledge, you might then want to encourage the team members to exchange ideas with a certain frequency, or divide their research time between individual learning (e.g., doing lab work on their own) and social learning (e.g., having meetings with colleagues) in a certain proportion, or put in place a certain hierarchy based on particular qualifications of the personnel, and so on.

The question of how best to combine individual and social learning is a broad one. To address it, we build on two existing, already well-understood frameworks in social epistemology, which will automatically limit our studies to specific ways of combining individual and social learning.

One framework is the so-called bounded confidence (BC) model, as developed in Hegselmann and Krause (2002, 2006, 2009, 2015, 2019). While in the original version of that model individual learning is treated as a black box, leaving the details of how agents update on the evidence they receive unspecified, we follow Douven and Wenmackers (2017) and Douven (2019, 2022a, 2022b) in opening the black box and modeling individual learning via broadly Bayesian update rules.

The BC model features communities whose members learn both individually and by attending to their peers' opinions. The model does not account for our reliance on *expert* opinion, however, or more generally, for social learning from successful members of one's social community. This form of learning is at the center of Schurz' (2019) recent work on meta-induction, which is the other framework to be used in our paper. In Schurz' work, different forms of learning are studied in the form of so-called prediction games, which take place in the form of consecutive rounds $i = 1, 2, \dots$, where in each round i a new data item d_i is presented, and

future data items, or alternatively their probabilities, are predicted. Schurz' prediction games involve communities consisting of *individual learners*, who adapt their beliefs solely on the basis of incoming pieces of evidence, together with *meta-learners*, who constantly monitor the beliefs and success records of individual learners and other meta-learners, basing their own beliefs on the beliefs of them. Particularly important for the present study are *inductive* individual learners, so-called object-inductivists, who base their epistemic method on some rule of induction from evidence, and inductive meta-learners, so-called meta-inductivists, who base their beliefs on the beliefs of other agents in proportion to their observed accuracy or closeness. We also look at agents who combine, in a variety of ways, object- and meta-inductivist learning.

In the epistemic framework of Schurz (2008, 2019), a given epistemic agent is confronted with a set of accessible epistemic methods (each method represented by an agent). The question is which of these methods, or combination of these methods, the agent should choose for the given cognitive task, which in Schurz' setting is a prediction task. Based on results in machine learning (Cesa-Bianchi & Lugosi, 2006), Schurz shows that there is a particular kind of meta-inductive method that combines the predictions of the accessible methods in proportion to their observed success and that is *universally optimal* in the following sense: In all possible environments, its predictive long-run success is guaranteed to be at least as high (and often higher) as (than) the best accessible method, with small upper bounds for short run losses (so-called regrets) that quickly vanish when the number of rounds increases. This result holds even in "adversarial" environments in which the hitherto best methods worsen their performance as soon as the meta-inductivist puts weight on them, as well as in irregular environments in which the success rates do not converge and the hitherto best (accessible) method or methods permanently change.

The form of meta-induction studied in this paper differs from Schurz' in two respects. First, in Schurz' setting, the meta-inductivist learns from all agents or methods that are accessible to her, averaging over their opinions or predictions weighted by their observed predictive success. In contrast, in this paper meta-induction is implemented in a social setting in the spirit of Hegselmann and Krause's BC model, where several meta-inductivists can mutually learn from each other.

Second, in Schurz' setting the meta-inductivists always combine the *actual* predictions of the accessible agents, that is, those delivered in the present round. This presupposes that the meta-inductivist delivers its prediction *after* the other agents (the meta-inductivists' "candidate methods") have delivered their predictions. In a social setting, in which meta-inductive agents mutually learn from each other, this is impossible and would lead into a "learning circle," since each agent would have to wait for the other agent's opinion before she can make up her own opinion. To avoid this circle, the agents in the BC model base their social learning on the *past* opinions of the other agents, specifically, on their asserted opinion in the *previous* round. Because this study investigates meta-induction in a social setting within the context of the BC model, the investigation focuses on meta-induction of the BC type. The study of Schurz-type meta-induction in a social setting would require a departure from the BC framework that is left to further studies.

The difference between Schurz-type meta-induction and BC-type meta-induction is epistemologically significant, since for BC-type meta-induction, a general optimality theorem fails. This becomes clear when BC-type meta-induction is applied to oscillating event sequences. For example, if an expert predicts the sequence 0101 . . . perfectly but a BC-type meta-inductivist bases her predictions on the expert's prediction in the previous round, then the meta-inductivist

will always predict incorrectly. However, in typical BC scenarios a *constant* parameter τ is predicted or estimated. In this case, the difference between Schurz-type meta-induction and BC-type meta-induction may be small, provided the amount of new evidence that enters the actual round is small or insignificant.

In what follows, by “meta-induction” we always mean meta-induction of the BC type. Because in a Schurz-type framework there are no mutual epistemic interactions between the meta-inductivists and the other agents, those models cannot serve very well to study social learning in general.

In the context of social learning, there is a clear opportunity to improve both on the BC model and on the Schurz models. As said, in reality agents in a community of inquirers learn in a variety of ways, most notably: (i) by exploring the world themselves, as the individual learners in Schurz’ models as well as the agents in the BC model do; (ii) by mutually exchanging ideas and beliefs with their peers, as the agents in the BC model, but not those in Schurz’ models, do; and (iii) by heeding the advice of those agents they deem to be experts (i.e., highly successful), as the meta-inductivists in Schurz’ models, but not the agents in the BC model, do. The obvious approach to bringing these forms of learning together into one model is to implement the idea of meta-induction in a framework with epistemically interacting agents, in the manner of the BC model.

That is the approach to be taken in this paper. There are many different routes to integrating the various forms of learning in a model. We proceed one step at a time, gradually building up the fully integrated model, paying special attention to the introduction of meta-inductivists in the BC model, which is done here for the first time. The main questions to be addressed along the way are *normative* ones, asking which mixes of forms of learning serve our pursuit of truth best. We evaluate the learning strategies to be considered with respect to their ability to make our beliefs more accurate, or more truth-like. Moreover, given that human abilities for learning, including social learning, are plausibly thought to have resulted from a process of natural selection (e.g., Boyd & Richerson, 1985; Henrich & Boyd, 1998, 2002; Boyd, Richerson, & Henrich, 2011; Richerson, 2019), we also evaluate learning strategies with respect to their evolutionary sustainability, reckoning with the possibility that some mixes of individual and social learning may bring short-term benefits, at least to some, but be long-term detrimental. Note that, in the settings to be considered in this paper, survival will most plausibly be interpreted in a metaphorical sense, for instance, in terms of postdocs having/not having their temporary contracts renewed, senior staff members being/not being able to hire former students or acquire further grant money, and so on.

Within the frameworks chosen for our study, the more specific questions our overarching research interest gives rise to include the following: Supposing the object-inductivists equipped with peer learning are given access to the meta-inductivists’ beliefs, will they benefit from that? Moreover, are there any benefits to meta-inductivists if at the same time they also, to some extent, explore the world directly on their own? Do they benefit from interacting with other agents in the way the agents in the BC model interact? Might communities of agents perform better, both at an individual level and at the group level, if all their members engage both in object-inductivist learning and in meta-inductivist learning, possibly at different times?

We use computational agent-based models to study these and other questions, and more broadly to compare different combinations of the various ways of learning mentioned in the above. To address the issue of sustainability, we use a standard evolutionary algorithm.

2 Theoretical background

We start by introducing the main elements to be relied on in our attempt to develop a model of learning that integrates individual learning with different forms of social learning.

2.1 Carnapian inductive reasoners

We assume, as intimated, a broadly Bayesian background, meaning that the belief states of the agents to populate our models are characterized by probability assignments to whichever hypotheses are of interest. More specifically, we assume a particular implementation of Bayesian inductive reasoning, as formalized through Carnap's (1952) λ rules. In the following, these rules serve as the core of any procedure of updating the agent's belief or prediction conditional on evidence coming directly from the world. That kind of updating is going to be contrasted, but also combined, with specific forms of social updating, where the latter refers to an updating operation that is at least partly based on the belief states of certain (possibly all) other agents in the community.

Throughout the following, different forms of learning will be studied in the form of *epistemic games* (which is a generalization of Schurz' prediction games). Epistemic games take place in consecutive rounds $i = 1, 2, \dots$, where in each round a new data item is presented, and the unknown value of some unknown parameter is estimated or predicted. In the epistemic games studied in this paper, the data items are the outcomes of a discrete random variable and the predicted parameter is their probability, Pr . Given elementary outcomes $K = \{Q_1, \dots, Q_k\}$ and given a sequence of n outcomes $e_n \in K^n$, where each Q_i has a fixed probability p_i of occurring, such that $p_i \geq 0$ for all i and $\sum p_i = 1$, let n_{Q_i} designate the number of outcomes in e_n in which Q_i was realized. Then Carnap's inductive learning rule is as follows:

$$\text{Pr}(Q_i | e_n) = \frac{n_{Q_i} + \frac{\lambda}{k}}{n + \lambda}. \quad (\text{C})$$

In all our computer experiments, the data will always concern a binary random variable (think, for instance, of coin tosses, the coin having a fixed but at least initially unknown bias), so that we in effect only look at instances of the above schema in which $k = 2$.

Note that (C) provides a continuum of inductive rules, yielding a different rule for different values of the λ parameter. This parameter regulates the learning rate: larger values will make learning from data slower, by making the effect of the agent's prior probabilities felt for a longer period of time; conversely, smaller values of λ mean that the data will more quickly wash out the priors, the limiting case being $\lambda = 0$, which means that the updated probabilities will always equal relative frequencies. So, for instance, if an agent sets λ to 0, then that agent's probability for a coin landing heads, given that so far 100 tosses with that coin produced 57 heads, equals .57. Carnap thought of λ as expressing an agent's degree of epistemic caution, which he considered to be a subjective matter.

2.2 Carnapian peers: The bounded confidence model

Instead of looking at Carnapian agents in isolation, we will study communities of such agents, where these agents can epistemically interact with others in their community. In modeling the epistemic interactions, we follow the lead of Hegselmann and Krause (2002, 2006, 2009), who present an agent-based computational model in which agents change their opinions partly by

“averaging” (in some way) the opinions of those epistemically close (in some sense) to them and partly by being sensitive to evidence they receive directly from the world. The aim of the BC model is to cover, in an idealized way, some fundamental aspects of our epistemic situation, which is characterized by the fact that we learn both by listening to others and by exploring the world with our own senses. The BC model has been widely used to investigate various descriptive issues, such as the conditions under which communities of initially disagreeing agents come to an agreement and those under which we should instead expect polarization to occur (e.g., Lorenz, 2003, 2008). However, increasingly the model is also used to address normative questions, for instance, concerning the practice of assertion (Olsson, 2008), the resolution of disagreement amongst peers (Douven, 2010; De Langhe, 2013), efficient truth approximation (Douven & Kelp, 2011; Douven, 2019), and practices of protecting us against mis- and disinformation campaigns (Douven & Hegselmann, 2021).

The agents in the BC model are all interested in determining the value of some parameter $\tau \in [0, 1]$. Each agent proceeds by updating its estimate of τ at discrete points in time (corresponding to rounds of the game) on the basis of (i) information about τ it receives from the world, and (ii) the estimates of τ of the agents within its bounded confidence interval (BCI), meaning that their estimate of τ is within some (typically small) distance of the agent’s own estimate.¹ Formally, agent x_i ’s opinion concerning τ after the $(n + 1)$ -st update (i.e., in round $n + 1$) is defined to be

$$x_i(n + 1) = \frac{1 - \alpha_i}{|X_i(n)|} \sum_{j \in X_i(n)} x_j(n) + \alpha_i \tau, \quad (\text{BC})$$

with $x_j(n)$ being the opinion of agent x_j after update n , and

$$X_i(n) := \{j : |x_i(n) - x_j(n)| \leq \epsilon_i\}$$

is the set of agents within agent x_i ’s BCI after update n , where $\alpha_i \in [0, 1]$ is a parameter determining the weight the agent gives to the “evidential” part of the updating relative to the “social” part, and $\epsilon_i \in [0, 1]$ is a parameter determining the agent’s BCI (which has a length of 2ϵ). As the subscripts indicate, agents can have their own α and ϵ values, but it is also possible to let all agents share those values.

As Hegselmann (2004) already noted, the BC model is highly flexible and can easily be tweaked to more specific research needs. And indeed, by now various extensions of the model exist. Among these are the extensions proposed in Douven (2010), Douven and Riegler (2010), and De Langhe (2013), which are all aimed at investigating situations in which agents receive “noisy” evidence; Crosscombe and Lawry’s (2016) model with agents whose opinions are allowed to be interval-valued; the extensions found in Lorenz (2003, 2008), Jacobmeier (2004), and Riegler and Douven (2009), in which agents have beliefs about more than one issue at the same time rather than about just one parameter; and Douven and Hegselmann’s (2021) model, which extends the original BC model by introducing typed agents (the types differing by the extent to which they act in epistemically responsible ways) as well as the mechanism of confidence dynamics, which allows agents to change their α and ϵ values in the course of the updating process.

All the aforementioned extensions of the BC model share with the core model Hegselmann and Krause’s explicit design decision to treat the updating on worldly evidence as a kind of black box: while there is a clear impact of worldly evidence on agents’ beliefs, there is no *explicit*

¹Agents are referred to as “it,” given that we will only be dealing with computational agents.

rule the agents are following to accommodate that evidence. Douven and Wenmackers (2017) present an extension that opens the black box. Specifically, in their version of the BC model, an agent’s belief state at any given time is characterized by a probability function defined on a set of self-consistent, mutually exclusive, and jointly exhaustive hypotheses and the updating on worldly evidence is unpacked in terms of Bayes’ rule and a probabilistic version of the so-called Inference to the Best Explanation (Douven, 2022b). The aim of Douven and Wenmackers’ (2017) study is to compare those rules with each other along various epistemically relevant dimensions (such as their accuracy; see below).

Douven (2022a) offers a somewhat different variant of the BC model that also opens the black box, equipping some agents with an instance of Carnap’s object-inductive rule (C), making others randomizers (in that, at each update, they randomly pick a new opinion), and again others dogmatists (in that they randomly pick an opinion at the start and then stick with that opinion throughout the whole updating process). The main simulations reported in Douven (2022a) focus on the first, “Carnapian” type of agents. In what follows, we call this type of agents “Carnapian peers,” or “Carnapians” for short. In the present paper, randomizers and dogmatists will play no role. (We do look at different types of agents as well, which will be introduced shortly.)

In Douven’s model, all agents receive binomial data (i.e., independent and identically distributed binary data), thought of as the outcomes of coin tosses, presented to the agents one toss after the other. In the simulations he ran, all agents received data from their own individual coin, where however all coins in any given simulation had the same bias or probability $\tau \in [0, 1]$ for heads, which was determined randomly at the beginning of the simulation. The value of τ was initially unknown to the agent, although it did know that all coins involved—its own coin and those of the other agents in the community—had the same bias. Each agent started with its own random estimate of τ . The agents updated that estimate after each toss, forming a new opinion, which for the Carnapians occurred as follows. Where e_n^i designates the sequence of outcomes that agent i has seen after the n -th update, and with H designating heads and $\Pr^i(H | e_n^i)$ being agent i ’s estimate of τ (and thus the probability of obtaining heads on a toss) after having seen data e_n^i , then for all Carnapians i and updates n ,

$$\Pr^i(H | e_{n+1}^i) = \frac{1 - \alpha_i}{|X_i(n)|} \sum_{j \in X_i(n)} \Pr^j(H | e_n^j) + \alpha_i \frac{(n+1)_H + \frac{\lambda_i}{2}}{n+1 + \lambda_i}, \quad (\text{BCC})$$

with

$$X_i(n) := \{j : |\Pr^j(H | e_n^j) - \Pr^i(S | e_n^i)| \leq \epsilon_i\}$$

the set of agents within agent i ’s BCI after the n -th update. Less formally, the estimated τ value of a given Carnapian in round $n+1$ is an α -weighted average of its τ value conditional on the evidence as given by the λ rule, and the average of the estimated τ values of all Carnapians in the previous round whose estimate was ϵ -close to that of the given Carnapian.

Figure 1 illustrates the kind of effects that changing the parameter setting can have on how the updating process unfolds. For instance, the top left graph shows the updating of a community of 50 Carnapians which do not epistemically interact at all, while the top right graph does the same for Carnapians which do give a lot of weight to the opinions of others in their community. We see in particular that, even though in the latter case the agents are quite conservative in regarding an agent worthy of impacting their own opinion (given the small value for ϵ), the fact that there is this interaction at all offers tremendous help speeding up

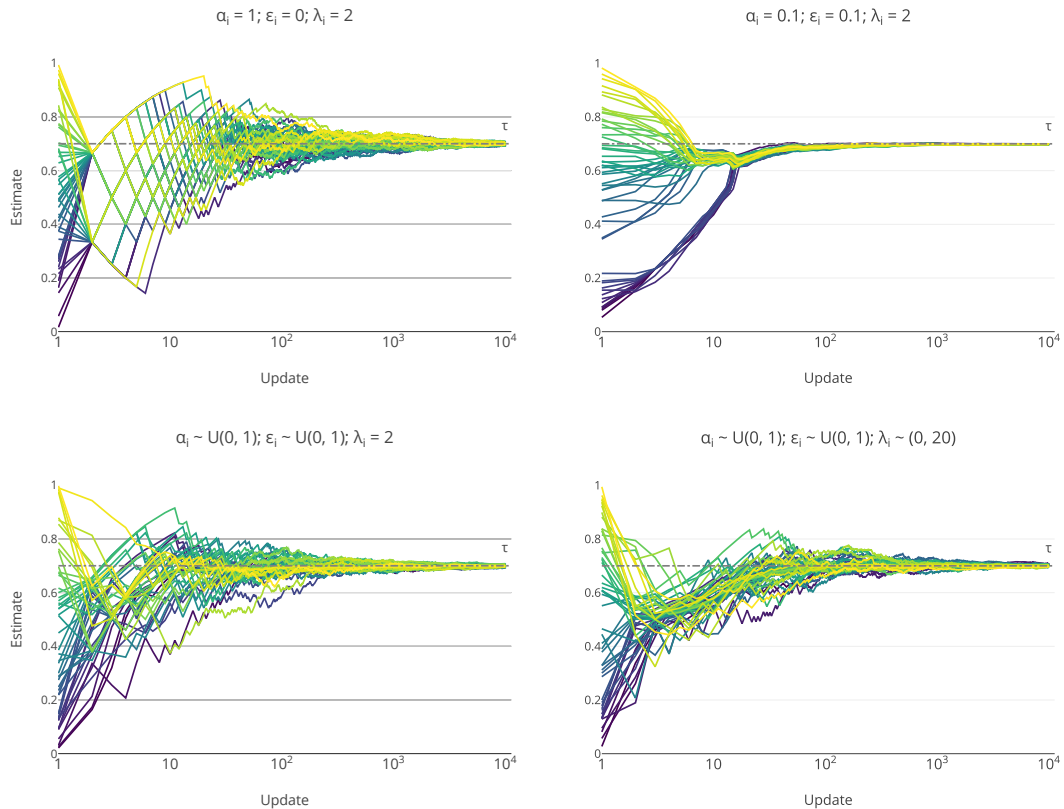


Figure 1: Log plots of repeated BC updating of communities of 50 Carnapians, for $\tau = 0.7$ with different parameter settings. (To facilitate tracking individual agents, agents are colored according to their initial estimates, with lighter colors corresponding to higher estimates.)

convergence of opinions. Comparing the graphs on the bottom row of the same figure, it is clear that we can expect an effect on speed of convergence of the agents' λ values as well.

2.3 Scoring

A concept that will play an important role throughout the following is that of an agent's score, which is meant to formalize the intuitive idea that some predictions are more accurate than others. An agent assigning probability p to heads at a given point in time can be interpreted as predicting a $p \times 100$ percent chance of heads on the next toss. If the next toss comes up heads indeed, we will deem the agent more accurate the closer p is to 1. If the agent predicted heads with certainty ($p = 1$), we will say it was entirely accurate. If the agent predicted, say, a 70 percent chance of heads, we will say that it was more accurate than if it had predicted a 50 percent chance, or a 30 percent chance, or indeed any percentage below 70. While there is an ongoing debate about how to formalize the notion of accuracy (see, e.g., Selten, 1998; Bickel, 2007; Douven, 2022b), a widely used measure of inaccuracy is the so-called Brier scoring rule (Brier, 1950), which for the case at hand equals $(1 - p)^2$ if the toss came up heads, and p^2 if it came up tails (with, as before, p being the agent's probability for heads). We assume this metric in all of the following. Note that Brier scores are penalties: lower is better as that indicates greater accuracy. We will thus mostly refer to the scores as *losses*.

2.4 Meta-induction

The machinery introduced so far was already explored in Douven (2022a) in the context of explaining the success of inductive reasoning. That work was meant as a complement to Schurz' (2019) work on meta-induction, which addresses an age-old philosophical question, to wit, that of how to justify our reliance on induction. Hume (1748/2006) had famously argued that induction—making predictions on the basis of past observations—is unjustifiable, however routinely we engage in that form of reasoning. For—as he noted—it is not a priori that induction is reliable, meaning that a *deductive* justification is not forthcoming; nor can we justify induction *inductively*, as that would involve us in circular reasoning. Over the centuries, various authors have attempted to find flaws in Hume's argument, but the consensus is that none of those attempts has been successful. The new observation that has fueled Schurz' response to “Hume's problem” is that we can rationally rely on induction in the absence of a proof that induction is reliable. The argument he develops is an *optimality* argument: whether or not induction is reliable, we can do no better than to rely on it. Drawing on results from machine learning, most notably the work by Cesa-Bianchi and Lugosi (2006) on prediction with expert advice, Schurz (2008, 2019) is able to give an analytic justification for meta-induction, by showing that, in every possible world, one can only do worse by not relying on meta-induction. More exactly, there is a particular kind of meta-inductive method that combines the predictions of the accessible methods in proportion to their so-called regrets and is universally optimal in the sense explained in Section 1.² Put briefly, the regret of a given agent is the difference between the total score of the agent and that of the meta-inductivist. The next step of Schurz' epistemological argument then consists of an application of meta-induction to the empirical findings about the predictive accuracy of our actual inductive practices as compared to various non-inductive methods. To the extent that certain object-inductive methods have been observed to be more successful than non-inductive methods in the past, we are meta-inductively justified to continue using these methods in the future. Because there is an independent optimality justification of meta-induction, this argument is no longer circular and in this way offers a solution to Hume's problem.

Douven (2022a) is particularly concerned with this second step: demonstrating the success profile of different kinds of inductive methods. Although Douven (2022a) was, in a clear sense, concerned with meta-induction, there were, among the agents comprising the societies that figured in the computer simulations reported in that paper, no meta-inductivists to be found. In this paper, we explore ways of combining the idea of meta-induction with the model of social updating presented in Douven (2022a).

2.5 Research questions

In the next section, we study the effects of individual learning and social learning in communities consisting of Carnapian peers and meta-inductivists (the former learning from the latter, and vice versa). For the reasons explained in Section 1, our meta-inductivists will be of the BC type, which means that they base their own opinion updates on the opinions the Carnapians held in the previous round, weighted by their regrets. A natural question to ask is

²In the machine learning literature, Hume's problem has been deepened in the form of Wolpert's (1996) no-free-lunch theorem (see also Wolpert & Macready, 1997), which roughly says that every possible prediction method has the same expected success averaged over all possible worlds. Schurz (2019, Sect. 9.3) proposes a novel solution to the no-free-lunch challenge based on the universal optimality of meta-induction.

how meta-inductivists would fare as members of the sort of society studied in Douven (2022a), and specifically how different types of meta-inductivists would fare. Furthermore, it would be interesting to know whether there might also be any benefits to the Carnapians—who do not attend to the past performance (Brier losses) of other agents but only to whether other agents’ opinions are close enough to their own—to be members of a community that includes some meta-inductivists. Also, in the standard BC model, as well as in the extension of that model defined in Douven (2022a), whether one agent’s opinion impacts another agent’s opinion depends strictly on the distance between their opinions (given a general ϵ value, or in the case the agents have their own ϵ value, on the second agent’s ϵ value).

In Section 4, we investigate an arguably more realistic alternative to the combination of individual learning and social learning as in Carnapian peers. We call this new form of mixed learning “merit-based learning,” and the respective agents “meritocrats.” Merit-based learning combines individual learning with social learning based on the other agents’ merits (or epistemic success) instead of on their closeness to the learner’s opinion. Thus, the social component of merit-based social learning is meta-inductive learning, but in merit-based learning meta-induction is combined with individual learning in the form of a weighted average.

In short, three types of epistemic methods will be studied: (i) Carnapian peers, (ii) meta-inductivists, and (iii) meritocrats. Furthermore, we look at forms of mixed learning where agents sometimes engage in merit-based updating, and sometimes explore the world entirely on their own, which is another way of bringing meta-induction into the framework of the BC model. In all cases, the overarching question will always concern the degree to which the various forms of learning are conducive to our epistemic goals. Like in Schurz’ work, we will be looking at the success of individual agents, but we are also interested in the success of groups as a whole, where the latter focuses in particular on whether the various strategies for social learning are evolutionary sustainable.

2.6 Methods

While computer simulations of agent-based models have long been used in the social sciences and in epidemiology, they are a more recent tool in social epistemology, which can be thought of as the study of social factors contributing to or impeding our quest for knowledge. Social epistemologists use agent-based models for broadly the same reason these models are being used elsewhere: they offer essential help studying phenomena that cannot be understood by focusing on agents outside their social context, or even by focusing on small groups of epistemically interacting agents (e.g., agents sharing their opinions with each other), and that are difficult or even impossible to investigate in a strictly analytical fashion. This also motivated the use of computational agent-based models in Douven (2022a), where it was argued that we have to look at the population level, and in particular at the evolutionary dynamics taking place at that level, to understand how humans became adept inductive reasoners.

In this paper, we take the same broad perspective, and therefore will also rely on computational agent-based models as well as on techniques of evolutionary computation. But the simulations will now focus not only on Carnapian peers in the context of the BC model but will also take meta-inductivists and meritocrats into consideration. By doing so, we hope to move closer to a complete model of social learning, more complete, at any rate, than is currently to be found in the literature.

3 Adding meta-inductivists

We start by considering communities of Carnapian peers to which we add one or more meta-inductivists. As already indicated, we work with a simple statistical model, in which each agent receives the outcomes of repeated tosses with the same coin. Although for each agent a separate coin is tossed, all coins have the same bias, which at least at start time is unknown to the agents. The Carnapians update their estimates of that bias as defined by (BCC). Note that (BCC) does not require the set of agents within an agent's BCI to consist of Carnapians only (or even of Carnapians at all). Accordingly, whenever in the following meta-inductivists are part of the community of agents, Carnapians take their opinions into consideration as well in the social part of the updating process. So, in particular, whenever a meta-inductivist is within a Carnapian's BCI, then the former's opinion weighs as heavily in the latter's updating as do the opinions of the Carnapians within its BCI.

All agents—both the Carnapians and the meta-inductivists—are scored after each update, in the way explained previously. Meta-inductivists update their opinions on the basis of regrets, which for every meta-inductivist and every Carnapian are computed as the difference between the meta-inductivist's loss and the Carnapian's loss. More exactly, after each update n , the meta-inductivists start by adding up, for each Carnapian, the Brier losses received by the Carnapian up until and including update n , and by doing the same for their own losses. Next, they calculate, again for each Carnapian, the difference between their own total Brier loss and the Carnapian's total Brier loss after update n . If their own total loss is bigger at the time, they assign the difference between the total losses as a weight to the Carnapian after update n ; else, the Carnapian receives a weight of 0 after that update. Provided any Carnapians receive a weight greater than 0, meta-inductivists make their new estimate of the bias (and hence their probabilistic prediction for the $n + 1$ -st piece of evidence) a weighted average of the Carnapians' estimates after the n -th update, the weights being determined as explained. If all Carnapians have a weight of 0, the meta-inductivist sets its new bias equal to that of the best-performing Carnapian (i.e., its probabilistic prediction for the next coin toss is whatever the best-performing Carnapian predicts; if there is more than one best-performing Carnapian, ties are resolved in some conventional manner, e.g., by picking the first in some given ordering of the Carnapians).

To give an impression of how the opinions of the agents in such mixed communities evolve, Figure 2 shows log plots of communities consisting of 50 Carnapians and 10 meta-inductivists, sequentially updating 10,000 times. In these illustrations, the Carnapians receive as worldly evidence tosses of coins with a bias of 0.7, those in the community shown in the left panel having α and ϵ values both of 0.1 and a λ value of 2, those in the community shown in the right panel having α and ϵ values drawn randomly and uniformly per agent from the (0, 1) interval and λ values drawn randomly and uniformly per agent from the (0, 20) interval.

The first question to be addressed is whether the Carnapians benefit from the presence of meta-inductivists in the community. We cannot expect the question to have a general answer, as Carnapians can vastly differ from one another in terms of α , ϵ , and λ values. We therefore tackled the question by running simulations that systematically went through an interesting range of values for one of the parameters while keeping fixed (in particular ways) the constraints on the other two. Needless to say, one could go even more systematically through parameter space here, though at considerable computational cost, at least if the results are to reach some level of reliability.




Figure 2: Log plots of repeated BC updating of communities of 50 Carnapians and 10 meta-inductivists, for $\tau = 0.7$ with different parameter settings. (Note that the parameters only pertain to the Carnapians.)

Specifically, we conducted simulations featuring communities all of which consisted of a fixed number of 50 Carnapians and with a variable number of meta-inductivists, ranging from 0 to 50, added to them. All agents started with a random estimate of the value of τ , where it was given that this value was in the $(0, 1)$ interval. For each of the α , ϵ , and λ parameters, we fixed constraints on two of them and then went in small incremental steps through a range of values for the remaining one, which was always the same for all Carnapians. For each specific setting, we let all Carnapian agents update by dint of (BCC) on the outcomes of 100 coin tosses that they received one per time step and where the coins all had the same bias that was randomly chosen at the start of each simulation. The meta-inductivists (if present) also updated at each time step, in the way just described. For each setting, we ran 100 simulations, after which we calculated the average total Brier loss over all 100 updates, the average being both over the Carnapian agents and over the 100 simulations. The constraints on the parameters not at issue were of two kinds: either the parameters had some fixed setting, or they were drawn randomly at the beginning of each simulation, in the manner indicated in the plot titles of the separate panels of Figure 3, which summarizes the results of these simulations.

To clarify further, each panel in Figure 3 shows a 51×51 grid, with the x coordinate of a cell indicating the value of either the α , or the ϵ , or the λ parameter and the y coordinate indicating the number of meta-inductivists present in the community. The color of each cell represents the average total Brier loss incurred by the Carnapians, where this average is itself averaged over the 100 simulations, each starting by randomly choosing a value for τ .

There are a number of observations that we can immediately make. First, we see that, both for the fixed parameter setting and for the setting with parameters drawn within bounds, increases in the α and ϵ parameter lead to lower average total Brier losses, although it appears that most benefits come from increasing these parameters from 0 to some still rather small value, and that further increases have little to no additional effect. Second, for the λ parameter the opposite appears to hold: smaller values are clearly associated with lower average total Brier losses. Third, and most importantly for present concerns, there appears to be little to no effect of increasing the number of meta-inductivists in the community on the average total Brier losses of the Carnapians. The top right corner of the left graph on the middle row, corresponding




Figure 3: Total Brier loss for the Carnapians, averaged over 100 simulations and over the 50 Carnapians in each simulation. See the text for further explanation.

to higher ϵ values and higher numbers of meta-inductivists, seems slightly darker than the bottom right corner of that graph, but that is the only effect that visual inspection suggests.

A statistical analysis is more revealing. To assess the effect of adding meta-inductivists, we can, for instance, in the simulations in which we systematically varied the value for ϵ , compare the average total Brier loss incurred by the Carnapians when no meta-inductivists were present with those incurred when one meta-inductivist was present and also with those incurred when 50 meta-inductivists were present. When we do this for each value of $\epsilon \in \{0, 0.01, 0.02, \dots, 0.5\}$,³ then a one-sample t -test shows that the mean difference in scores for the case with 0 meta-inductivists and that with 1 meta-inductivist is not reliably different from 0, neither for the setting with $\alpha = 0.1$ and $\lambda = 2$ nor for that with $\alpha_i \sim \mathcal{U}(0, 1)$ and $\lambda_i \sim \mathcal{U}(0, 20)$; $t(50) = 1.49, p = .14$ and $t(50) = 1.52, p = .13$, respectively. By contrast, making the same comparisons for the case with 0 meta-inductivists and that with 50 meta-inductivists shows that the mean difference in total average Brier losses is reliably greater than 0, for both parameter settings; $t(50) = 10.11, p < .0001$, Cohen's $d = 1.41$ (which is a large effect size) and $t(50) = 5.19, p < .0001$, Cohen's $d = 0.73$ (medium), respectively. This indicates that the presence of 50 meta-inductivists helped to reliably reduce the Brier losses of the Carnapians, in both kinds of situations we considered.

Figure 4 gives more detailed insight in the relation between average total Brier losses and the presence of 50 meta-inductivists in the case in which the effect of the latter's presence is most outspoken (i.e., the case with $\alpha = 0.1$ and $\lambda = 2$). This figure shows the dependence of the help the presence of meta-inductivists can offer on the value of ϵ , which—recall—represents how liberal the Carnapians are in counting others as their peers, whose opinions they want to take into account. In the figure, Δ represents the difference in average total Brier score incurred by the Carnapians in communities without meta-inductivists and that incurred by the Carnapians in communities with 50 meta-inductivists. We see a clear upward trend in which higher ϵ values tend to be associated with higher values for Δ , meaning that as ϵ goes up, the tempering effect of the presence of the meta-inductivists on the Carnapians' losses gets larger.

The same comparisons for the simulations in which we systematically varied the value of α yielded similar results for the setting with $\epsilon_i \sim \mathcal{U}(0, 1)$ and $\lambda_i \sim \mathcal{U}(0, 20)$. For this, we found no reliable difference in losses when 1 meta-inductivist was added but a highly reliable one when 50 meta-inductivists were added; $t(50) = 0.01, p = .99$ and $t(50) = 6.05, p < .0001$, Cohen's $d = 0.85$ (large), respectively. All results for the other parameter setting were non-significant. The pattern was the same for the remaining set of simulations, in which we varied the value of λ : for the setting with $\alpha_i \sim \mathcal{U}(0, 1)$ and $\epsilon_i \sim \mathcal{U}(0, 1)$, the presence of the 50 meta-inductivists significantly lowered the average total Brier losses of the Carnapians, as compared to the situation without any meta-inductivists ($t(50) = 5.56, p < .0001$, Cohen's $d = 0.78$), but none of the other comparisons yielded a significant result.

Whereas one could explore parameter space more systematically and more exhaustively than was done in the above, the results so far already allow us to answer one of our research questions: provided there are *enough* meta-inductivists present in a community, their presence can have a lowering effect on the Carnapians' Brier losses. Whether it has, depends on which values the Carnapians have adopted for the α , ϵ , and λ parameters. This means that, while it may appear that, and may in fact be the case that, the meta-inductivists are exploiting the Carnapians as a kind of epistemic slaves, who are doing the hard work of collecting and assessing

³Note that, here, all agents always have the same ϵ value. So, ϵ without the subscript i is to be read as "for all i , $\epsilon_i = \dots$ "; similarly for α and λ further on.




Figure 4: Differences between average total Brier loss incurred by 50 Carnapians in communities without meta-inductivists and average total Brier loss incurred by the Carnapians in communities with 50 meta-inductivists (Δ), for values of $\epsilon \in \{0, 0, 01, 0.02, \dots, 0.5\}$. For all Carnapians, $\alpha = 0.1$ and $\lambda = 2$. See the text for further explanation. (Regression line added to highlight the trend; shaded area showing 95 percent confidence interval.)

evidence, it turns out that, on the supposition that the Carnapians are also *social* learners, the Carnapians still get something in return for being exploited.

We have so far looked only at the performance of the Carnapians. How are the meta-inductivists doing in terms of accuracy? Might they not be better off by becoming Carnapian learners? We ran further simulations, similar in their set-up to the ones just reported, but now always with next to 50 Carnapians also 50 meta-inductivists present in the community. In particular, every simulation started by picking a random value for $\tau \in (0, 1)$ and also picking random start opinions for each of 50 Carnapians and each of 50 meta-inductivists. In all simulations, the agents repeatedly updated 100 times, the Carnapians on the outcomes of 100 coin tosses that they received sequentially, all tosses coming from coins with a bias determined by the value of τ picked for the given simulation, and the meta-inductivists by basing their opinions on those of the Carnapians, in the manner previously explained. We ran 100 of such simulations for two parameter settings, which determined the α , ϵ , and λ values for the Carnapians, to wit, one with $\alpha = \epsilon = 0.1$ and $\lambda = 2$, and one with $\alpha_i \sim \mathcal{U}(0, 1)$, $\epsilon_i \sim \mathcal{U}(0, 1)$, and $\lambda_i \sim \mathcal{U}(0, 20)$. In all simulations, we measured Brier losses both for the Carnapians and for the meta-inductivists.

Figure 5 shows box plots of the outcomes of the two sets of simulations. It appears that, on average, the meta-inductivists incurred lower Brier losses than the Carnapians. The difference was significant for both sets of simulations, as confirmed by two two-sample t -tests. For the setting with $\alpha = 0.1$, $\epsilon = 0.1$, and $\lambda = 2$, in which Carnapians incurred an average total Brier loss of $1.11 (\pm 0.62)$ and meta-inductivists an average total Brier loss of $0.38 (\pm 0.20)$, we obtained $t(198) = 10.58$, $p < .0001$, Cohen's $d = 1.48$ (large). For the setting with $\alpha_i \sim \mathcal{U}(0, 1)$, $\epsilon_i \sim \mathcal{U}(0, 1)$, and $\lambda_i \sim \mathcal{U}(0, 20)$, in which Carnapians incurred an average total Brier loss of $0.82 (\pm 0.55)$ and meta-inductivists an average total Brier loss of $0.52 (\pm 0.30)$, the results of the t -test were similar: $t(198) = 4.98$, $p < .0001$, Cohen's $d = 0.7$ (medium).

The findings so far suggest that, while meta-inductivists do better on the count of accuracy, Carnapians benefit as well from being members of a mixed community. However, there are questions of fairness to be addressed. Given how we set up things, we can easily picture the Carnapians as being the only ones to get their hands dirty, and so one might expect them




Figure 5: Box plots showing the outcomes of 100 simulations for two different parameter settings, each measuring the average total Brier loss incurred by 50 Carnapians and that of 50 meta-inductivists together forming a community.

to complain about the fact that the meta-inductivists do better than they (the Carnapians) do while living off the latter's honest toil. But even meta-inductivists should be concerned about the sustainability of the arrangement. While meta-inductivists do better from a short-term perspective (looking only at single simulations), if accuracy determines agents' chances of evolutionary success (i.e., chances of survival and of having successful offspring), it is by no means obvious that the kind of community we are considering here favors the survival of meta-inductivists as a special sort of agent.

To examine this, we recruited an evolutionary algorithm, which models computationally the mechanisms of variation and selection (Coello Coello, 1999). In detail, the evolutionary computations consisted of the following steps:

1. We started with a community of 100 agents, 50 of which were Carnapians, and 50 meta-inductivists. For the Carnapians, α , ϵ , and λ values were picked randomly, with $\alpha \sim \mathcal{U}(0, 1)$, $\epsilon \sim \mathcal{U}(0, 1)$, and $\lambda_i \sim \mathcal{U}(0, 20)$. The initial opinions of both the Carnapians and the meta-inductivists were chosen randomly and uniformly from the unit interval. The initial Brier losses were (of course) 0.
2. Each agent received a piece of evidence at each time step, where this piece of evidence could again be thought of as the outcome of a coin flip carried out at that time. This was done individually per agent, so that we may think of each agent as having its own coin which was flipped at each time step. It was given that all coins had the same bias, and it was the agents' task to estimate this bias.
3. At each time step, the agents updated their opinions on the evidence they had received at that time, Carnapians using (BCC) with their specific parameter values plugged in, and meta-inductivists in the manner described previously in this section.
4. Also at each time step, each agent was scored, using the Brier rule (see Sect. 2.3), on the basis of the evidence obtained at that time. For instance, if at the previous time step the agent had updated its opinion to .6 (for heads) and the newly produced evidence was heads, then we would add a loss of $(1 - .6)^2$ to the agent's score. Thus, an agent's score at any given time was its cumulative Brier loss, the sum of the Brier losses incurred up and till that point.

5. A single run started by picking a value for τ randomly and uniformly from the unit interval. Then, for each agent, a coin with bias τ was flipped 100 times, where at each of those times the agent would update its opinion about the bias on the evidence obtained at that time. At the end of the run, the sum of the agent's losses, over all 100 updates, was calculated.
6. For each generation, 25 such single runs were conducted, after which, for each agent in that generation, its average total Brier loss (averaged over the 25 runs) was calculated.
7. Then 50 agents were selected on the basis of their average total Brier loss. Specifically, the reciprocals of those averages, normalized so as to sum to 1, served as survival probabilities, that is, as the probability for the agent to be among the 50 that were to form the parent population of the new generation. They formed this new generation together with 50 children. Each selected Carnapian C generated offspring by partnering with a random other selected Carnapian C^* and producing a Carnapian child whose α -, ϵ -, and λ -parameters were a mix of those of C and C^* , in that the values were chosen randomly and uniformly from the intervals spanned by the parents' values for the relevant parameters; the child's starting opinion was picked randomly and uniformly from the unit interval. Lacking α -, ϵ -, and λ -parameters, selected meta-inductivists procreated by producing a child with a random starting opinion chosen in the same way. Both parents and children entered the new generation with their scores set (for parents, rather, reset) to 0.
8. This was repeated for 50 generations, after which the process was terminated.

We repeated this procedure 100 times.

Of these 100 simulations, 43 terminated early, on average after 14.44 (± 24.92) generations, simply because the Carnapians had vanished completely from the community. In the other 57 simulations, the communities in the last generations consisted predominantly of meta-inductivists, containing on average only 30.61 (± 11.94) Carnapians. Thus, if evolutionary pressure comes strictly from inaccuracy, then choosing to be a meta-inductivist, in the way considered here, is a risky bet. After all, the success of this kind of meta-inductivism can be self-undermining, and rather short-lived. In the simulations that terminated early, the meta-inductivists' seeming success led to the quick demise of the Carnapians, and once the last Carnapian was gone, the meta-inductivists were doomed as well. Needless to say, there are ways to prevent this from happening. The most obvious way is to furnish meta-inductivists with an object-inductive fallback method, which is to kick in if there are no object-inductivists around whose estimates can be imitated or combined. All meta-inductive methods studied in Schurz (2019) are equipped with such a fallback method that is employed if there are no agents with positive weights. As fallback method we prefer a Carnapian method of learning, identical to that used by Carnapian peers. We call meta-inductivists with such a fallback method "self-sustained meta-inductivists," as opposed to "pure meta-inductivists" that are doomed to die when there are no more individual learners around. Re-running the evolutionary simulations of this section with self-sustained meta-inductivists leads to the effect that those simulations that earlier on terminated early would now continue with a community of Carnapian peers. Since this result is not particularly interesting, we abstain from presenting this variant of the simulations.




Figure 6: Log plots of repeated BC updating of communities of 50 Carnapians and 10 meritocrats, for $\tau = 0.7$ with different parameter settings. (Note that the ϵ -values only pertain to the Carnapians.)

4 Adding meritocrats

In Section 2.5, we indicated that a potentially fruitful way to combine meta- and object-inductivism in the BC model is to make the updating of some agents in that model dependent not on the distances among opinions but upon agents’ losses. For any agent of that type, the *social* part of the updating process would come to consist of taking a weighted average of the opinions of all other agents in its community, the weights being based on those other agents’ losses.

Like the meta-inductivists that we examined in the previous section, the new type of agent also weighs the Carnapians’ opinions on the basis of those agents’ merit; whether the Carnapians’ opinions are close to the agent is immaterial to the latter. We call the new agents “meritocrats.”

Schurz (2019) understands meta-induction as a research program within which a variety of different meta-inductive strategies are investigated. Viewed from this perspective, meritocrats are a sort of meta-inductivists, since their social learning component is based on meta-induction. Unlike the earlier meta-inductivists, however, meritocrats attend to the evidence they receive from the world (which will again be coin tosses). They also have an α and a λ value, which have the same interpretation as they have for Carnapians (i.e., determining how to weigh worldly evidence against social updating, and determining attachment to prior opinion, respectively). Then, where the index i ranges over all and only meritocrats and the index j over all and only Carnapians, the meritocratic updating procedure comes to the following, for all updates n :

$$\Pr^i(H | e_{n+1}^i) = \frac{1 - \alpha_i}{\sum_j w_n^j} \sum_j w_n^j \Pr^j(H | e_n^j) + \alpha_i \frac{(n+1)_H + \frac{\lambda_i}{2}}{n+1 + \lambda_i}, \quad (\text{BCM})$$

with w_n^j being equal to 1 divided by the total Brier loss incurred by Carnapian j up until and including the n -th update. We furnish meritocrats with a fallback method right away. The method is simply to let the index j in (BCM) range over all and any agents present in the community if no Carnapians are around.

As we did for meta-inductivists, we give an impression of meritocratic updating by showing in Figure 6 the repeated updating of a mixed community, consisting of 50 Carnapians and 10 meritocrats, for the same parameter settings that were considered in Figure 2. Note that, for

the setting in which α equals 0.1 for all agents, after one update the meritocrats may hold one of two opinions: all meritocrats who received heads as outcome will share the same opinion, and so will all meritocrats who received tails as outcome. After the second update, we may find four different opinions among the meritocrats, then eight, and then ten (because there are only ten meritocrats). This is somewhat visible in the left panel of Figure 6. For the other setting, of course, in which all meritocrats may have their own α value, we may find them all to have different opinions at any time.

We can raise all the same questions that were raised in the previous section, now with respect to meritocrats. In particular, do the Carnapians benefit from having meritocrats around them? How do Carnapians and meritocrats compare on the count of accuracy? Is it a wise strategy, from an evolutionary standpoint, to be a meritocrat, supposing selection occurs on the basis of accumulated Brier losses?

We addressed the question of whether Carnapians benefit from the presence of meritocrats in the same manner in which we addressed the question of whether they benefit from the presence of meta-inductivists. We reran the simulations whose outcomes we visualized in Figure 3, the only difference being that now the communities consisted of 50 Carnapians and a variable number of meritocrats instead of Carnapians and meta-inductivists. The results were similar to those obtained in the simulations with meta-inductivists, as can already be seen (somewhat) from comparing Figure 7 with Figure 3. In the former figure, as in the latter, cell color represents the total Brier loss incurred by the Carnapians, where this total is averaged over the 50 Carnapians in each simulation and then over the 100 simulations.

The statistical analysis also gives more or less the same results as for the earlier simulations: In particular, adding one meritocrat to a community otherwise consisting of Carnapians never helped to significantly reduce the Brier losses of the latter. Adding 50 meritocrats, instead, did reliably reduce Brier losses in the simulations, for the simulations in which we varied ϵ for both settings, and for the remaining simulations only for the random settings (all $ps < .0001$). So here, too, we can conclude that, under circumstances, the presence of enough meritocrats in a community helps the Carnapians to become more accurate. We can also look again at how the two types of agents compare in terms of overall accuracy. The results indicate that, depending on the parameter setting, meritocrats can do significantly better than Carnapians. Figure 8 shows the outcomes for two different parameter settings. Running two-sample t -tests for these results shows that, whereas in the random setting the difference in average Brier losses is not significantly different between the two groups ($t(198) = 1.67, p = .09$), the difference is highly significant for the other setting ($t(198) = 7.82, p < .0001$, Cohen's $d = 1.10$), in which the Carnapians incur an average total Brier loss of $1.08 (\pm 0.66)$ and the meritocrats one of $0.52 (\pm 0.29)$.

We recruited the same evolutionary algorithm we used for looking at the sustainability of the meta-inductivists' update strategy to look at how communities consisting of Carnapians and meritocrats would evolve. We found that in 100 simulations starting with a community of 50 Carnapians and 50 meritocrats, the Carnapians vanished completely in 65 of them. By contrast, the meritocrats vanished completely in only 11 of the simulations. On average, there were $19.96 (\pm 35.29)$ Carnapians in the final generation. Note that for meritocrats the question of sustainability does not arise, as they already have a fallback method.

The success of the meritocrats reveals an interesting truth about the BC model, viz., that rather than grounding the social part of the updating in similarity of opinion, we should ground it in merit. In other words, it makes a lot of sense to devote the same amount of research, if not




Figure 7: Total Brier loss for the Carnapians, averaged over 100 simulations and over the 50 Carnapians in each simulation. See the text for further explanation.




Figure 8: Box plots showing the outcomes of 100 simulations for two different parameter settings, each measuring the average total Brier loss incurred by 50 Carnapians and that of 50 meritocrats together forming a community.

more, that has been devoted to the original BC model also to the redefined model that replaces similarity-based social updating by merit-based social updating. To forestall misunderstanding, there is in our view a place for both a similarity-based and a merit-based model. For even if, from a normative standpoint, merit-based social updating is preferable, it is undeniable that similarity-based social updating does occur (Sunstein, 2019). And for various reasons—for instance, for those explained in Douven and Hegselmann (2021)—it is valuable to have a model that allows us to study that kind of updating computationally.

5 All together now

We have compared Carnapian peers with two types of broadly meta-inductive agents performing meta-induction over the Carnapians: pure meta-inductivists (of the BC type), and meritocrats that combine meta-induction with individual learning. The meta-inductivists were equipped with Carnapian learning as fallback strategy, as explained in Section 3. An obvious further step is to look at communities featuring all three types of agents. Figure 9 illustrates what an updating process in such a mixed community may amount to.

Running 100 simulations with similarly mixed communities, now consisting of 50 agents of each type, in which at the start of the simulation a value for τ is randomly set and then all agents update—in the way specific to their type—on the outcomes of 100 consecutive tosses with their own private coin with bias τ , we registered again the total average Brier losses for the three types of agents (so, per simulation averaged over the 50 agents belonging to each of the types). This was done for the two parameter settings we have been using for previous illustrations. They are indicated by the panel titles in Figure 10, which visualizes the results.

An ANOVA for the first parameter setting revealed that agent type had a significant effect on accuracy: $F(2, 297) = 64.97, p < .0001$, with an associated ω^2 of 0.30, which counts as a large effect size. Pairwise comparisons showed that Carnapians (with an average Brier loss of $0.95, \pm 0.58$) were significantly ($p < .0001$) less accurate than both meta-inductivists ($0.38, \pm 0.25$) and meritocrats ($0.42, \pm 0.26$), but that meta-inductivists and meritocrats did not differ significantly in terms of accuracy. For the other parameter setting, we found a significant effect of agent type on accuracy as well: $F(2, 297) = 20.51, p < .0001$, with $\omega^2 = 0.12$, which




Figure 9: Log plots of repeated BC updating of communities of 50 Carnapian agents, 10 meta-inductivists, and 10 meritocrats, for $\tau = 0.7$ with different parameter settings. (Note that the α and λ values do not pertain to the meta-inductivists and the ϵ values only pertain to the Carnapians.)

is a medium to large effect size. Post hoc tests showed that the mean average Brier loss of meta-inductivists ($0.51, \pm 31$) was significantly ($p < .0001$) lower than that of Carnapians ($0.93, \pm 0.61$) and also significantly ($p < .0001$) lower than that of meritocrats ($0.83, \pm 0.53$). The means of Carnapians and meritocrats did not differ significantly from each other.

That meta-inductivists came out best here was not too surprising, given what was already observed in the previous sections. The results from the evolutionary computations were only slightly more exciting. The setup was essentially the same as that of the evolutionary computations reported in the previous section, the main differences being that now the community always also included 50 meritocrats, next to 50 Carnapians and 50 meta-inductivists, and that each simulation ran for 100 generations, unless in an earlier generation two of the three groups of agents had vanished, in which case the simulation was terminated. The α and λ values of children of selected meritocrats were picked in the same way as the α , ϵ , and λ values of the children of selected Carnapians were picked (i.e., randomly from the interval spanned by the parents' corresponding values; see the previous section). In these simulations, Carnapians vir-




Figure 10: For 100 simulations with communities of 50 Carnapians, 50 meta-inductivists, and 50 meritocrats, average total Brier scores for each of the groups, for two different parameter settings.

tually never made it until the end; there were on average only $4.5 (\pm 21.4)$ Carnapians to be found in the last generations. Recall that we endowed the meta-inductivists with a fallback method to be used in the absence of Carnapians. So whenever the Carnapians vanished in a simulation, the meta-inductivists became a weighted average of individual learners and social peer learners, forming their opinion in the same way the Carnapians had done, which in this case meant that they took into account the opinions not only of their fellow meta-inductivists but also of the meritocrats. That not only allowed the meta-inductivists to survive, but also to win, even if not by a lot: in the last generations, there were on average $85.16 (\pm 63.39)$ meta-inductivists versus an average of $60.34 (\pm 63.71)$ meritocrats. That difference is significant ($p < .01$) but the effect size is small ($d = 0.39$).

6 Undisclosed group membership

So far, the agents always knew what type another agent belonged to. Because of that, meta-inductivists were able to update strictly on the evolving opinions of the Carnapians and meritocrats were able to base the social part of their updating on those opinions only. But the assumption that agents wear their type on their sleeve is far from realistic. One may wonder what difference it makes if this assumption is given up, so that the social updating (which is *all* the updating, in the case of meta-inductivists) always happens on the basis of *all* the agents' opinions.⁴ Formally, Carnapians keep updating via (BCC) while meritocrats also keep updating via (BCM), albeit that now the index j in these equations ranges over *all* agents in the community, and not just over the Carnapians. For meta-inductivists, now, their estimate of τ after n updates—so, their probability for heads on the $n + 1$ -st toss—is the weighted average of the estimates of τ after n updates of *all* agents in the community, where the definition of weights remains as before.

We repeated all the same simulations from the previous section, now under the assumption that agents did not know to which group their fellow-agents belonged, and so assuming the update mechanisms as just explained. Figure 11 shows the outcomes from 100 simulations in which we calculated the average Brier losses for the three groups of agents, each 50 strong and jointly forming a community. ANOVAs showed that in the setting with fixed parameters, type had a significant and large effect on accuracy: $F(2, 297) = 23.27, p < .0001, \omega^2 = 0.13$. For the other setting, the results were not significant. For the former setting, Carnapians (with a mean total Brier loss of $1.04, \pm 0.61$) were significantly less accurate than meta-inductivists ($0.63, \pm 0.51$) as well as meritocrats ($0.57, \pm 0.45$), at an α level of $.0001$, but meta-inductivists and meritocrats did not differ significantly from each other.

Because the scores were so close here, it was not immediately clear what to expect from the evolutionary computations. Here, these computations got interesting indeed, so we state the algorithm we used again in some detail:

1. We started with a community of 50 Carnapians, 50 meta-inductivists, and 50 meritocrats. For each Carnapian, α , ϵ , and λ values were chosen randomly at start time. The first two parameters were again chosen uniformly from the $(0, 1)$ interval and the third from the

⁴Note that for Carnapians this was true from the beginning. One could also consider the kind of situation in which for them social updating is done strictly on the opinions of the Carnapians within their BCI, but in that kind of situation it does not make sense to ask whether there might be any advantage for Carnapians to be members of a community that also includes meta-inductivists or meritocrats.

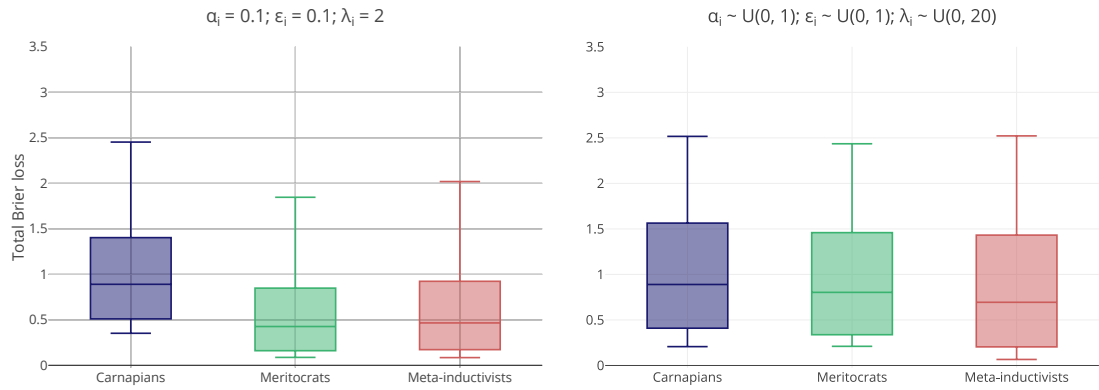


Figure 11: For 100 simulations with communities of 50 Carnapians, 50 meta-inductivists, and 50 meritocrats, average total Brier scores for each of the groups, for two different parameter settings.

(0, 20) interval. For each meritocrat, α and λ values were chosen in the exact same way (meritocrats do not have an ϵ value).

2. As previously, the agents were tested in single runs in which each of them received 100 pieces of evidence at discrete time steps; here too, every piece of evidence could be thought of as the outcome of a coin flip. The evidence was generated for each agent separately. It was again given that all coins had the same bias, and the task was also as before: the agents had to estimate the bias shared by all coins.
3. The agents updated on each piece of evidence they received in the way determined by their type and (for Carnapians and meritocrats) by their parameters. For meta-inductivists, this meant that they ignored the evidence and took a weighted average of all the agents' opinions about the bias, the weights being as defined previously.
4. Also as before, the evidence served to score an agent's opinion. The agents' cumulative Brier losses served as input for determining the weights needed for the meta-inductivists' and meritocrats' update rules (which, note, are defined differently, although in both cases being loss-based).
5. Single runs were exactly as in previous evolutionary computations: a value for τ was picked randomly and uniformly from the unit interval; for every agent, an initial opinion was picked also randomly and uniformly from the unit interval; and each agent's total score for that run was set to 0. The run ended by recording the sum of the agent's losses over the 100 updates.
6. Each generation went through 25 single runs, after which the total average Brier losses of all the agents in that generation were calculated.
7. Then 75 agents were selected, indeterministically, their survival probabilities again being based on their total average Brier losses, in the manner explained previously. Here, too, these 75 agents were the parent population for the next generation, which they formed together with 75 children that were created by first randomly drawing pairs from the parent population, where the draws were with replacement, but with the restriction that each pair had to consist of members belonging to the same type. Then, for Carnapians, a child was created by drawing α , ϵ , and λ values randomly and uniformly from the interval spanned by the two parents' α , ϵ , and λ values, respectively, which then became the child's α , ϵ , and λ values. For meritocrats, the procedure was exactly the same, except

that (of course) no ϵ value needed to be drawn. And for any pair of meta-inductivists (a type of agent without any parameter values), one child was created.

8. This was terminated if only one type of agents remained—in which case we say that the simulation *converged*—or after 100 generations.

We ran this procedure 100 times. Figure 12 shows, for six randomly picked simulations of those 100, the number of agents of each type present in the consecutive generations. We see that four of the six shown processes converged, in that one type took over the whole community. Of the 100 simulations, 34 did not converge within 100 generations. Figure 13 shows a histogram of the numbers of generations it took to convergence, where the ones represented as having taken 100 generations did in fact *not* converge.

The meritocrats were the clear winners. There were on average 108.16 (± 64.54) of them in the last generations. The other two groups did about equally poorly, with on average 22.12 (± 39.58) Carnapians present in the last generations and 19.72 (± 39.36) meta-inductivists. How to make sense of the fact that the meta-inductivists, which often came out best in the direct comparisons in terms of accuracy, do not do well in these simulations? The plots in Figure 12 already give a clue. In all of them, meta-inductivists do best for the first couple of generations. That is no coincidence: for a short while, meta-inductivists did best in *all* of the simulations we ran; see also Figure 14, which shows the average number of meta-inductivists present in a generation, the average being over the 100 simulations. However, soon the meritocrats start to catch up, and then tend to take the lead. To see why this happens, note that, while meritocrats can improve as a group, in that the evolutionary process can help fine-tune their parameters to values that are optimal for the kind of test situation we have placed the agents in (more on this below), this is not true for the meta-inductivists, which do not have parameters to be fine-tuned, and so which also do not really evolve as a group. To be sure, the Carnapians have parameters that can be fine-tuned as well, even one over and above the parameters the meritocrats have. These will be fine-tuned, too, in the process, which—as Figure 12 shows—sometimes keeps them in the running all the way till the 100-th generation (at which point the evolutionary algorithm was terminated). But that in general they lose to the meritocrats is for the reason we have already seen, to wit, that meritocratic updating, in particular the meta-inductive part of that updating which takes agents' past performance into account, is a better idea (for reasons of accuracy) than an updating rule which, for the social part, only looks at similarity of opinion.

We have explained the superiority of the meritocrats over the meta-inductivists in terms of the evolutionary fine-tuning of the meritocrats' parameters based on incoming evidence. Against this explanation it may be objected that also the meta-inductivists indirectly profit from this fine-tuning because they follow in their opinions the opinions of the most successful meritocrats. Recall, however, that the meta-inductivists studied in this paper are of the BC type and not of the Schurz type, which means that they base their opinions on the opinions of the meritocrats (and Carnapians) in the previous round, which reflect only past data but not the most recent data items. We conjecture that this constitutes a second reason for the evolutionary superiority of the meritocrats over the meta-inductivists. Especially in earlier rounds of a simulation run, the fact that meta-inductivists do not have indirect access to the most recent data incurs additional losses, compared to meritocrats with well-chosen parameters.



Figure 12: Counts of agent type per generation for six randomly chosen simulations. See the text for further explanation.

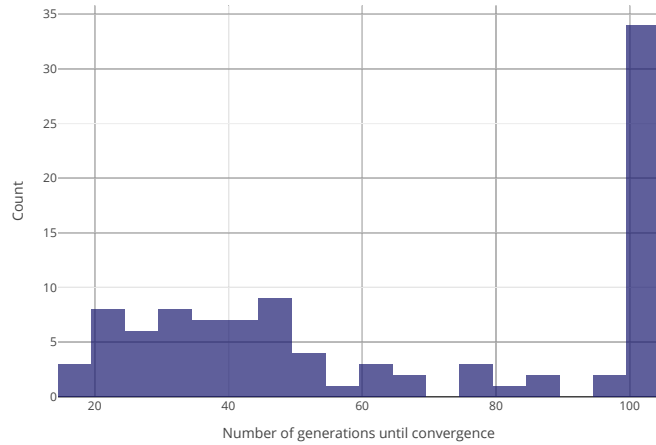


Figure 13: Histogram of number of generations it took in the 100 evolutionary computations to reach convergence. The peak entirely to the right indicates simulations that in fact did not converge.

7 Mixed learners

Previous sections considered agents that either (i) combine individual and social learning at any time step, where however the two components can be differently weighted by different agents, or (ii) only learn by looking at the opinions of others. Meta-inductivists made up the latter group, while in the former group, there were agents which also took the past performance of their fellow agents into account in updating, thereby being a kind of *partial* meta-inductivists. The group of agents that combined individual and social updating at all times also included the Carnapians, which for the social part looked at whether other agents were close enough to them, not on how those other agents had done so far.




Figure 14: Mean number of meta-inductivists present in a generation, averaged over 100 simulations; the shaded area indicates one standard error from the mean.

There is something clearly realistic about the simulations we ran. In reality, we all combine, in some way, individual and social learning. Also, we are probably all somewhat different in how we go about learning; as far as the social aspect is concerned, some people may, at least some of the time, just want to fit in with their neighbors, while others may be more inclined to seek expert advice. At the same time, there is also something quite *unrealistic* about the simulations. After all, rather than engage *simultaneously* in two sorts of learning—object-learning and meta-learning—we more typically engage in different types of learning at different times. More often than not, we have a choice to investigate a matter by looking for evidence in the world or to consult others, and we sometimes do the one, sometimes the other. Indeed, for any more complicated matter that may require some time to investigate, we may today do some research on our own and adjust our opinion based on the outcome of that research, and tomorrow talk to a number of colleagues and be influenced by their opinions. In a stylized form, that is what the agents to be considered in the current section do. Whether, at any given time, they are more likely to engage in research on their own or to seek the advice of their colleagues is determined by some “trait,” which we here formally model by means of a threshold value. In the simulations to be reported in the following, at each time step, and for each agent separately, a random value is drawn from the unit interval. If the value is above the agent’s personal threshold value, the agent does its own research; otherwise, it listens to its colleagues instead, where the weights it gives to their opinions are the reciprocals of their Brier losses. So, for instance, a threshold of .8 means that, at any given time, the agent will act as a social learner with a probability of 80 percent, and look at, and update on, the worldly evidence produced at the time with a probability of 20 percent. We call the new type of agent that operates in this way a “mixed learner.”

Formally, at any given time t , a mixed learner x_i can be thought of as a tuple $\langle \lambda_i, \theta_i, o_i^t, l_i^t \rangle$ of four parameters: its value for the λ parameter in the Carnapian schema (BCC); a threshold $\theta \in [0, 1]$ determining the likelihood it is to act as a meta-learner at t and update by taking the loss-weighted average of the other agents in its community rather than update on the worldly evidence revealed to it at t ; its estimate, o , at the time; and its loss, l , at the time.

Figure 15 shows a log plot of a run over 10,000 time steps with a community of 50 mixed learners, where the λ values (which determine individual learning rates) are chosen randomly and uniformly from the (0, 20) interval, and the θ values (which determine how likely an agent is to engage in individual rather than social learning) are chosen randomly and uniformly from the unit interval.

The single run offers no more than a sanity check: it looks like there is nothing out of the ordinary here. The question we want to focus on is whether there are optimal values, or rather optimal combinations of values, for the two parameters involved. For the above run, parameter values were chosen randomly for each agent. But one could ask whether agents would be better off if, for instance, they had a low value for θ and so would spend more time investigating the world on their own than talking to colleagues. Or might the optimal value for one parameter depend on the value for the other? For instance, might agents with a low θ value be best off with a high λ value and agents with a high θ value with a low λ value, where perhaps there is no significant difference between the two combinations? We cannot answer these questions a priori. Instead, we turn again to evolutionary computations to see what combination or combinations of parameter settings communities of agents competing for survival and an opportunity to reproduce evolve to.

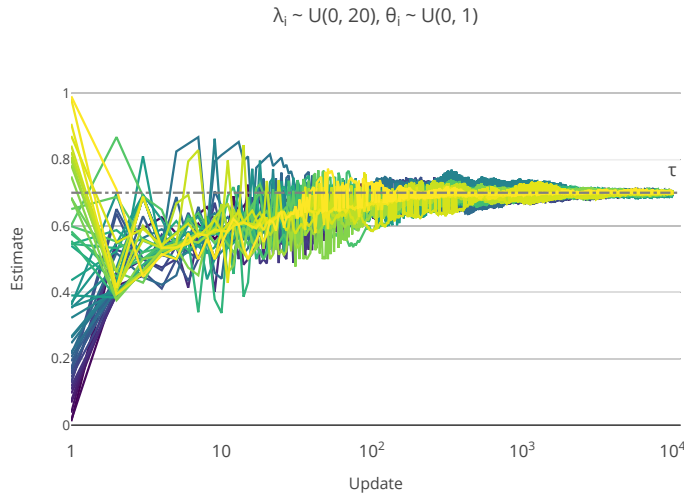


Figure 15: Log plot of a run over 10,000 time steps with a community of 50 mixed learners. (The same coloring convention is used as in Fig. 1.)

The evolutionary computations we ran proceeded along the lines of those reported in earlier sections:

1. We started with a community of 50 agents, whose λ and θ values were chosen randomly, specifically, $\lambda_i \sim \mathcal{U}(0, 20)$ and $\theta_i \sim \mathcal{U}(0, 1)$. The agents' initial opinions were chosen randomly and uniformly from the unit interval. Their initial scores were (of course) 0.
2. Each agent received a piece of evidence at each time step, where this piece of evidence could again be thought of as the outcome of a coin flip carried out at that time. This was done individually per agent, so that, as before, we may think of each agent as having its own coin which was flipped at each time step. Here, too, it was given that all coins had the same bias, and it was the agents' task to estimate this bias.
3. Whether the agent at a given time actually *looked* at the piece of evidence that, at this time, was produced for it, depended on whether a number randomly and uniformly drawn from the unit interval at the same time step exceeded the agent's personal threshold value. Only if it did would the agent look at the evidence and update its opinion about the bias on that evidence. Otherwise, it would update by taking a loss-weighted average of all the agents' opinions about the bias.
4. Whether or not the agent looked at the newly produced evidence, that evidence would serve to score the agent's opinion, that is, its estimate of the bias at the time the given piece of evidence was produced. An agent's loss at any given time was its cumulative Brier loss, the sum of the Brier losses incurred up and till that point.
5. A single run started by picking a value for τ randomly and uniformly from the unit interval. Then, for each agent, a coin with bias τ was flipped 100 times, where at each of those times the agent could either look at the result and update its opinion about the bias on it or ignore the result and update on the basis of the opinions of the other agents instead. At the end of the run, the agent's total loss was recorded.
6. For each generation, 25 such single runs were conducted, after which, for each agent in that generation, its average total Brier loss (averaged over the 25 runs) was calculated.




Figure 16: Per-generation box plots of the λ values, showing how they change during the evolutionary computations.

7. Then 25 agents were selected, in the same manner as in the evolutionary computations previously reported. These formed the parent population of the new generation, which further consisted of 25 children that were here created by taking random pairs from the parent population and drawing a λ value randomly and uniformly from the interval spanned by the pairs' λ values, as well as drawing a θ value randomly and uniformly from the interval spanned by the pairs' θ values. Those became their child's λ and θ values, respectively. The child's initial opinion was randomly and uniformly chosen from the unit interval and its initial loss was set to 0.
8. This was repeated for 20 generations, after which the process was terminated.

We repeated this procedure 100 times.

When we look at the results of the simulations, we see, in Figure 16, that the λ values converged rather quickly: from the seventh generation onward, λ values are, on average, below 5. In the last generation, the average λ value was 4.69. In Figure 17, we see a similarly quick convergence for the θ values, which on average are and remain a bit above .75 after four generations. In the last generation, the average θ value was .78.

We now have an indication of what, according to the outcomes of our evolutionary computations, is the optimal combination of parameter values. But how good is that combination, in terms of helping agents to reduce their overall Brier loss, compared with a combination of randomly chosen values, and, more importantly still, compared with forgoing social learning altogether and always only updating on the worldly evidence (so setting $\theta = 0$)?

To find out, we conducted further simulations. In these, we compared three different communities of agents in terms of their accuracy. One community consisted of agents that all used the best settings for the λ and θ parameters, as established by our evolutionary computations. That is to say, all agents had the same λ value of 4.69 and the same θ value of .78, which are the averages documented two paragraphs back. A second community instead used random λ and θ values, meaning that for each agent individually, a λ value was chosen randomly and uniformly from the (0, 20) interval and a θ value was chosen, also randomly and uniformly, from the unit interval. And the final community consisted of agents that were assigned a random value for λ

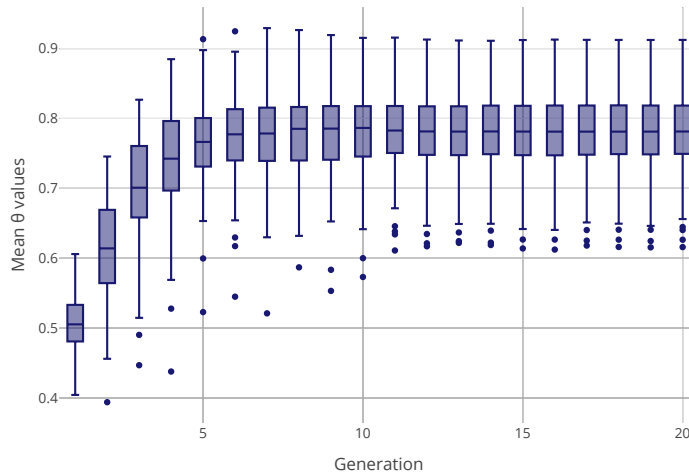


Figure 17: Per-generation box plots of the θ values, showing how they change during the evolutionary computations.

(determined in the same manner as for the previous community) but that all had a value for θ equal to 0, meaning that these agents *only* engaged in individual learning: at each time step, each agent updated on the piece of evidence that was offered to it at that time step and so never updated on the basis of the opinions of other agents in the community. Which (if any) of these settings is most advantageous, when it comes to accuracy?

We looked at this question again in the context of a coin-tossing model. Because we assumed that the answer could depend on the exact bias of the coin, we ran, for each possible bias in $\{0, .1, .2, .3, .4, .5\}$, 50 simulations for each of the three communities, where each separate simulation consisted of a single run of 1,000 updates. After each simulation, we calculated the total Brier loss incurred by the population (so the sum of the total Brier losses of all the agents in the population). The results are shown in Figure 18.

The pattern is more than clear—and quite remarkable. The community of nonsocial updaters appears to do, in general, much worse than the other two, and the community of agents with random values for λ and θ appears to do, across the board, quite a bit worse than the community using the best setting for those parameters. These impressions were confirmed by running an ANOVA per bias (smallest $F = 431.02$, largest $p < .0001$, smallest $\omega^2 = 0.85$). This is remarkable if one considers that, of the best-performing community, at each time step on average less than a quarter of the agents are attending to the evidence; all others are looking at the opinions held by the members of the community and are updating on the basis of those. There could hardly be a more convincing illustration of the benefits of social learning (as touted in, e.g., Goldman, 1999).

At the same time, these results raise a new question: We said that in the best-performing community less than a quarter of the agents are paying attention to the data, at any given point in time. But that will typically be a *different* fraction each time. Would one obtain the same results, or possibly even better ones, if there were a *fixed* division of cognitive labor in a community, with for instance 22 percent (corresponding to the best θ value of .78; see above) of the agents being designated to always attend to the data while the remaining agents update on the opinions of others? Simulation outcomes reported in Section 3 suggest that the resulting community,

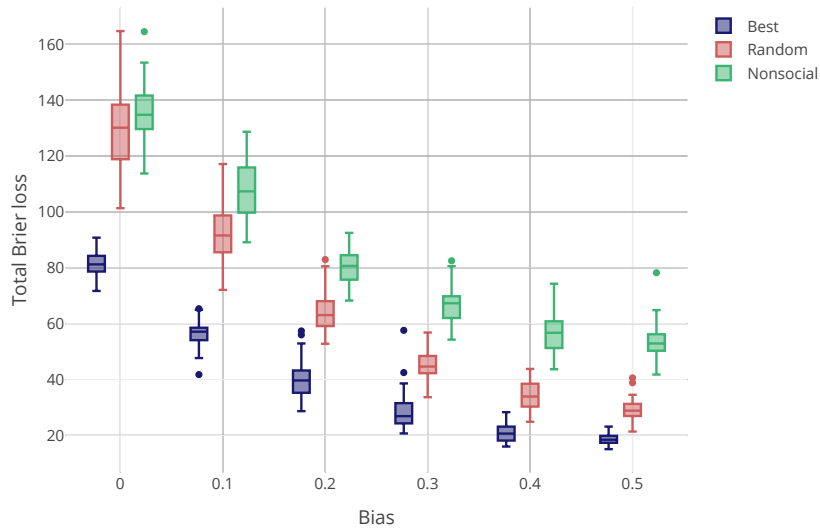


Figure 18: Box plots of total Brier losses over 1,000 updates for communities with three different settings for the λ and θ parameters: one community with the best setting (determined via evolutionary computing), one with random values for those parameters, and one with a “nonsocial” setting ($\theta = 0$ for all agents).

which would be a particular mix of object-inductivists and a kind of meta-inductivists, is not a sustainable entity. Naturally, there are ways to counteract evolutionary pressures. For example, if the community as a whole benefits from having among them a small percentage of object-inductivists, which would perish if evolution went unchecked, compensatory measures may be put in place to ensure that the object-inductivists get a fair share of the overall benefits. The meta-inductivists might profit from that as well. Of course, if no Carnapians are around, meta-inductivists will turn into Carnapians, by stipulation, but that should be a measure of last resort.

Again, whether this approach makes sense is hard to tell a priori. To shed light on the matter, we ran some further simulations, now comparing the previous best-performing community with the just-described mix of object- and meta-inductivists, more exactly, with a community of 50 agents, 12 of which were object-inductivists, the rest being meta-inductivists. Otherwise, the simulations were exactly the same as the previous ones.

The results are depicted in Figure 19. The mixed community does, in general, clearly much worse than the best-performing community from our previous experiment. And indeed, with the exception of the .4 and .5 biases, the best-performing community is significantly better, as confirmed by a series of t -tests ($p < .0001$, except for the .3 bias, for which $p < .01$).

In short, we wondered about optimal combinations of the relevant parameters and sought to determine those via evolutionary computing. We reckoned with the possibility that, for instance, a high value for θ and a low value for λ is about as beneficial for an agent as vice versa. But that is not what we found. The evolutionary process pushed the θ as well as the λ values of all agents into relatively narrow ranges. The most remarkable finding was that optimal θ values tended to be *high*, meaning that agents spent most of the time looking at other agents’ opinions at the expense of looking at the worldly evidence; in the terminology of March (1991) and Hills et al. (2015), they spent more time *exploiting* the information already

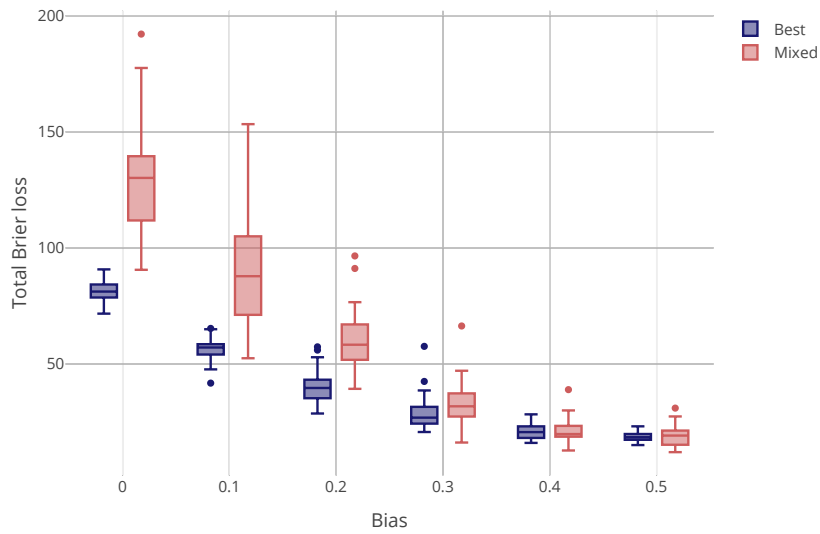


Figure 19: Box plots of total Brier losses over 1,000 updates for the best-performing community and a mixed community consisting of 12 object-inductivists and 38 meta-inductivists.

present in their community than they spent on *exploring* the world in order to contribute to the shared information. It was then further remarkable that we could not quite achieve the same results by assuming a fixed division of cognitive labor in a community of agents, with some agents only attending to the worldly evidence and others only attending at others' opinions. Apparently, to achieve optimal accuracy, it is essential that *all* agents spend some of their time investigating the world on their own, but then spend *most* of their time attending to the other agents' opinions, where—importantly—they also take into account those agents' track records.

As a kind of grand finale, we looked at how the new type of agents do in comparison with the types introduced in previous sections when they all have to compete for survival and procreation. We ran a further set of 100 evolutionary computations identical to those described in Section 6, except that now the starting generation also contained 50 mixed learners. With six exceptions, the computations ran for 100 generations, at which point they were terminated. The results from six randomly selected simulations shown in Figure 20 make it reasonable to believe that many of the ones that did not converge within 100 generations might never have converged, given that they seemed to have reached a sort of relatively stable equilibrium. From the figure, one guesses that meritocratic learning and mixed learning are both sustainable learning strategies, although the former appears to increase one's chances of evolutionary success more than the latter. Indeed, averaged over the 100 simulations, last generations featured 131.82 (± 41.61) meritocrats, 60.82 (± 33.09) mixed learners, and 7.36 (± 20.79) Carnapians. Meta-inductivists again never got very far.

8 General discussion

In the foregoing, we built on previous work extending the BC model, a popular model for studying communities of interacting agents. In particular, Douven (2022a) had already proposed to make the individual part of the updating mechanism explicit in terms of the Carnapian λ rules,

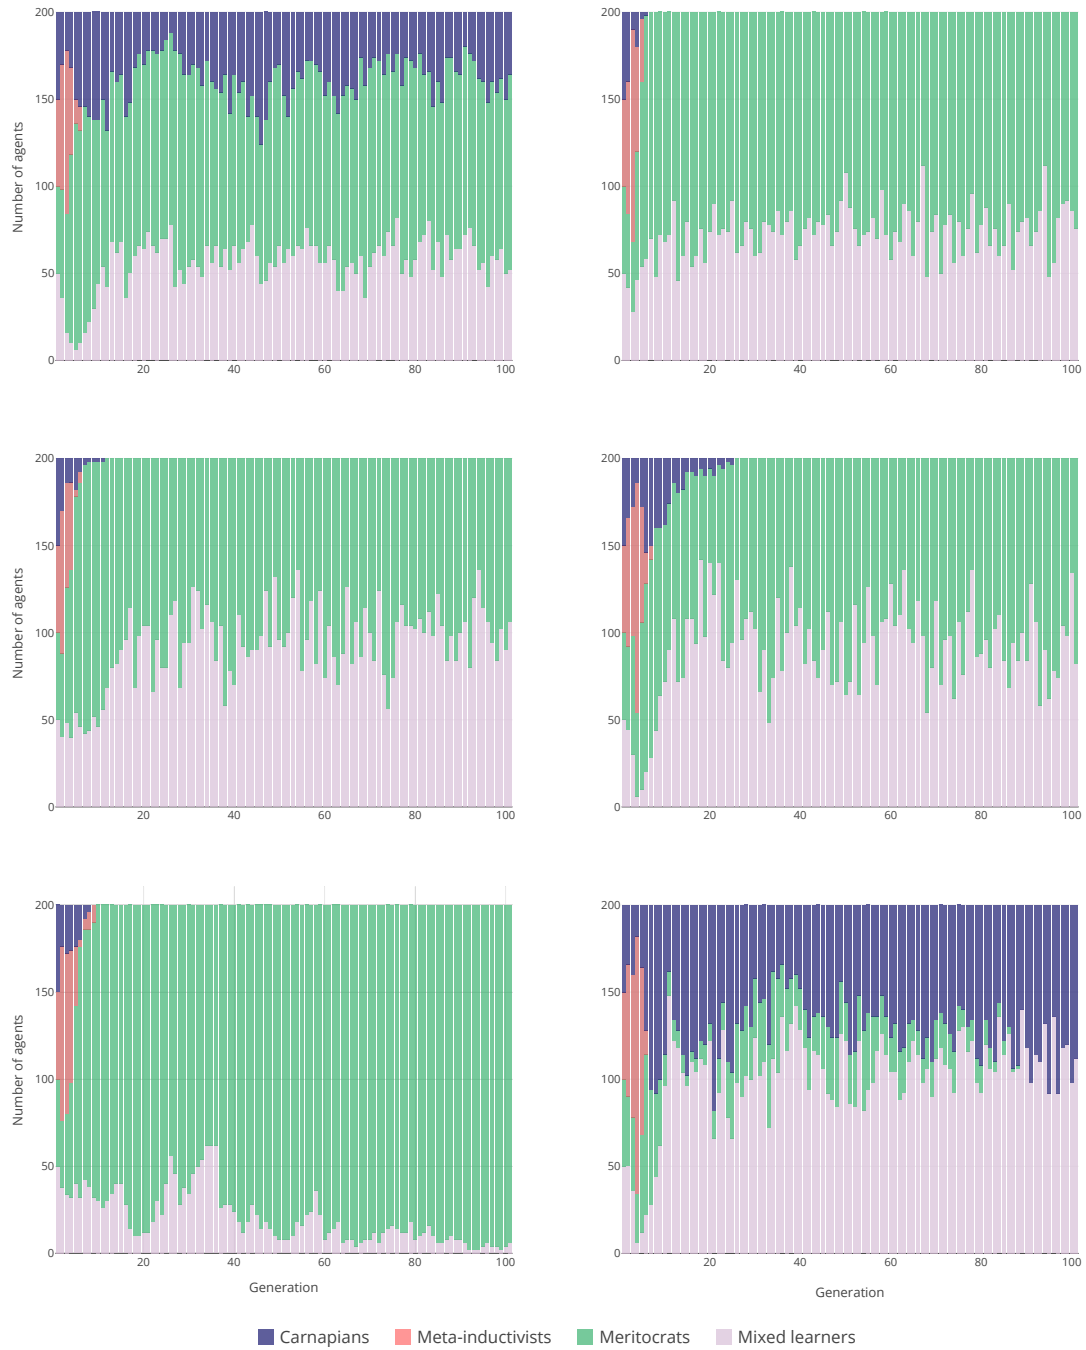


Figure 20: Counts of agent type per generation for six randomly chosen simulations. See the text for further explanation.

and Douven and Hegselmann's (2021) extension of the BC model already introduced the idea of working with communities consisting of differently typed agents.

Apart from the Carnapian peers, which were the only type of agents examined in Douven (2022a), the types studied in this paper were new in the context of the BC model, although they were all inspired by Schurz' (2019) work on meta-induction. First, there were the meta-inductivists in the "purest" form, modeled after Schurz' idea, but with the difference that their updating rule follows Hegselmann and Krause's BC-account, having access only to the agents' opinions in the previous round, not in the actual round. Their learning method consisted solely of merit-oriented social updating, which they carried out by taking weighted averages of the opinions of the Carnapians, the weights being a function of agents' Brier losses. Second, there were the meritocrats, which were a kind of conceptual amalgamation of Carnapians and meta-inductivists. Like Carnapians, they update on the worldly evidence provided to them, using some Carnapian λ rule, but also on the opinions of other agents, where however it is not (as in the original BC model) similarity of opinions that matters for this part of the updating procedure but rather past performance of the agents holding the opinions. Finally, there were the mixed learners, which sometimes act as meta-inductivists, sometimes as Carnapians (i.e., object-inductivists).

A general conclusion to be drawn from the foregoing is that meta-inductive reasoning is certainly a good idea, even if, probably, it is best *not* implemented in its purest form. The more realistic options of meritocratic learning and mixed learning also appear to be the better ones. The evolutionary computations we ran were stylized but do seem to model aspects of learning that *are* related to how well we do in our lives. At the same time, we do not believe that the model we used is realistic enough to license any definite conclusions about whether meritocratic or mixed learning is the more viable option. For instance, as already mentioned, it is unlikely that our evolutionary success depends entirely on the accuracy of our opinions.

This suggests an obvious avenue for future research, to wit, trying to make the evolutionary computations more realistic, most notably, by making success not only a matter of accuracy. Here, one could think of also looking at fast approximation to the truth, which can be in tension with accuracy (some learning strategies may take us relatively close to the truth very quickly but may then make it hard to get still closer to the truth or arrive at the exact truth; see, e.g., Douven, 2010), but which is certainly something we value in many practical contexts as well (Douven, 2022a, 2022b).

Also, we might be interested in strategies that protect us against mis- and disinformation campaigns, even if such strategies will sometimes slow us down in our efforts to get to the truth. In our simulations, all agents were always epistemically responsible in that they sought the truth. Douven and Hegselmann (2021) study models populated with agents some of which are epistemically *irresponsible* in that they try to deceive, for financial or political or other purposes, the agents which *are* epistemically responsible. Douven and Hegselmann show how different social learning strategies can help to mitigate, to varying extents, the harmful effects of the presence of deceivers. In Douven and Hegselmann (2021), the epistemically responsible agents are as defined in the original BC model, so with updating on worldly evidence treated as a black box mechanism. Also, their model involves no meta-inductive or meritocratic learning. It would be interesting to know whether those types of learning can further help to blunt the impact of mis- and disinformation campaigns. The most straightforward way to investigate this question would be to add epistemically irresponsible agents—for instance, "pseudo-Carnapians," who

pretend to update on evidence but in reality are pushing some falsehood—to the various models developed in this paper.

Another extension worth exploring adds networks to the models. Douven and Hegselmann (2022) distinguish between two kinds of networks: *access* networks and *impact* networks. For any given agent, the former determine the agents about whose belief states the agent is fully informed, while the latter determine the agents whose belief states impact the agent's belief state. As Douven and Hegselmann point out, access networks are entirely trivial in the standard BC model, given that every agent knows, or can be assumed to know, every other agent's belief state at any given time. Impact networks, by contrast, are dynamic: they can change from one update to another, given that the agents in a given agent's BCI can change. Douven and Hegselmann argue that the assumption of trivial access networks tacitly underlying the standard BC model is highly unrealistic and they therefore propose a new extension of the model with (static) *nontrivial* access networks added to it as a new layer. Exploring this extension, Douven and Hegselmann find network effects at the community level as well as at the individual level. Specifically, they find that more tightly connected communities tend to be more accurate, overall, and that individual agents tend to be more accurate the more *central* (as measured by standard centrality measures) their location is in the network. Douven and Hegselmann's extended model is populated strictly by Carnapians. It would be interesting to know what difference it would make to the various models studied in the present paper if they were made more realistic by adding (nontrivial) access networks to them. We intend to study this matter in future work.

Furthermore, our efforts to offer a comprehensive computational approach to studying social learning, combining individual learning with realistic forms of social learning, fall short in other respects as well. For one, our models feature agents whose belief states are exceedingly poor in comparison to real people's belief states. For another, individual learning is limited to inductive reasoning, in the form of probabilistic updating. To start with the latter, Schurz and Hertwig (2019, p. 19) argue that, whereas deductive reasoning has “maximum ecological validity,” by which they mean that it is a perfectly general form of reasoning that one can rely on in each and every situation, it has “very low applicability,” in the sense that “the prevalence of deductive inferences with *nontrivial* conclusions is low.” That—as they explain—is in contrast to non-deductive reasoning, which we in fact frequently rely on in our everyday lives, even though there are environments in which it can lead us astray. But while missing out on deductive reasoning might not be a great loss, it would be worthwhile to try and include abductive reasoning (roughly, reasoning on the basis of explanatory considerations) as well, although at the moment we lack good formal models for this type of reasoning (Douven, 2022b). A goal that is more easily achievable in the short term is to extend the models presented in this paper to ones that endow agents with richer belief states. Here, we could take our cue from the work of Lorenz (2003, 2008), Jacobmeier (2004), and Riegler and Douven (2009), among others.

Finally, despite the idealizations and limitations of the models described in this paper, we can raise the question to what degree these models are descriptively adequate. Even if our main goal was normative, this is an important question, given that, famously, “ought” implies “can.” In this connection, there are in fact two empirical questions (at least) that emerge from our findings. One is to what extent the various models approximate behavior, both at a collective and at an individual level, of groups of collaborating people, like researchers in a laboratory or employees of a company bent on achieving a common goal. The other is to what extent social engineers (e.g., lab managers) might be able to implement in non- or sub-optimally functioning

groups the kind of learning strategies that our findings would seem to recommend. In the introduction, we cited empirical research on how groups collaborate, and some of the methods used in that research could be instrumental in helping to answer the previously mentioned questions as well. We are hoping that these questions will be taken up by researchers better versed in the use of those methods than we are.

Acknowledgments

We are greatly indebted to three anonymous referees for valuable comments on a previous version of this paper.

Conflict of interest

The authors declare that they have no conflict of interest.

Data availability

Data sharing not applicable to this article as no data sets were generated or analyzed during the current study.

References

- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. (2017). Julia: A fresh approach to numerical computing. *SIAM Review*, 59, 65–98.
- Bickel, J. E. (2007). Some comparisons between quadratic, spherical, and logarithmic scoring rules. *Decision Analysis*, 4, 49–65.
- Boyd, R. & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Boyd, R., Richerson, P. J., & Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108, 10918–10925.
- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78, 1–3.
- Carnap, R. (1952). *The continuum of inductive methods*. Chicago: University of Chicago Press.
- Cesa-Bianchi, N. & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge: Cambridge University Press.
- Coello Coello, C. A. (1999). A comprehensive survey of evolutionary-based multi-objective techniques. *Knowledge and Information Systems*, 1, 269–308.
- Crosscombe, M. & Lawry, J. (2016). A model of multi-agent consensus for vague and uncertain beliefs. *Adaptive Behavior*, 24, 249–260.
- Darr, E. D., Argote, L., & Epple, D. (1995). The acquisition, transfer and depreciation of knowledge in service organizations: Productivity in franchises. *Management Science*, 41, 1750–1762.
- De Langhe, R. (2013). Peer disagreement under multiple epistemic constraints. *Synthese*, 190, 2547–2556.

- Douven, I. (2010). Simulating peer disagreements. *Studies in History and Philosophy of Science*, 41, 148–157.
- Douven, I. (2019). Optimizing group learning: An evolutionary computing approach. *Artificial Intelligence*, 275, 235–251.
- Douven, I. (2022a). Explaining the success of induction. *British Journal for the Philosophy of Science*, in press.
- Douven, I. (2022b). *The art of abduction*. Cambridge MA: MIT Press.
- Douven, I. & Hegselmann, R. (2021). Mis- and disinformation in a bounded confidence model. *Artificial Intelligence*, 291, 103415, <https://doi.org/10.1016/j.artint.2020.103415>.
- Douven, I. & Hegselmann, R. (2022). Network effects in a bounded confidence model. *Studies in History and Philosophy of Science*, 94, 56–71, <https://doi.org/10.1016/j.shpsa.2022.05.002>.
- Douven, I. & Kelp, C. (2011). Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis*, 75, 271–283.
- Douven, I. & Riegler, A. (2010). Extending the Hegselmann–Krause model I. *Logic Journal of the IGPL*, 18, 323–335.
- 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.
- Fortunato, S. (2004). The Krause–Hegselmann consensus model with discrete opinions. *International Journal of Modern Physics C*, 15, 1021–1029.
- Gaifman, H. (1986). A theory of higher order probabilities. In J. Halpern (ed.), *Theoretical aspects of reasoning about knowledge: Proceedings of the 1986 conference* (pp. 275–292). San Mateo CA: Morgan–Kaufmann.
- Goldman, A. I. (1999). *Knowledge in a social world*. Oxford: Oxford University Press.
- Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63, 85–110.
- Goldman A. I. (2010). Epistemic relativism and reasonable disagreement. In R. Feldman & T. A. Warfield (eds.), *Disagreement* (pp. 187–215). Oxford: Oxford University Press.
- Harris, P. & Corriveau, K. (2011). Young children's selective trust in informants. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366, 1179–1187.
- Hegselmann, R. (2004). Opinion dynamics: Insights by radically simplifying models. In D. Gillies (ed.), *Laws and models in science* (pp. 19–44). London: King's College Publications.
- Hegselmann, R. (2014). Bounded confidence, radical groups, and charismatic leaders. In F. Miguel, F. Amblard, J. Barceló, & M. Madella (eds.) *Social simulation conference advances in computational social science and social simulation*. Barcelona: Autonomous University of Barcelona, DDD repository, <http://ddd.uab.cat/record/125597>.
- Hegselmann, R. (2020). Polarization and radicalization in the bounded confidence model: A computer-aided speculation. In V. Buskens, R. Corten, & C. Snijders (eds.) *Advances in the sociology of trust and cooperation: Theory, experiment, and field studies* (pp. 197–226). Berlin: De Gruyter.
- Hegselmann, R., König, S., Kurz, S., Niemann, C., & Rambau, J. (2015). Optimal opinion control: The campaign problem. *Journal of Artificial Societies and Social Simulation*, 18, <http://jasss.soc.surrey.ac.uk/18/3/18.html>.

- 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. (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.
- Hegselmann, R. & Krause, U. (2015). Opinion dynamics under the influence of radical groups, charismatic leaders, and other constant signals: A simple unifying model. *Networks and Heterogeneous Media*, 10, 477–509.
- Hegselmann, R. & Krause, U. (2019). Consensus and fragmentation of opinions with a focus on bounded confidence. *American Mathematical Monthly*, 126, 700–716.
- Henrich, J. & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, 19, 215–241.
- Henrich, J. & Boyd, R. (2002). On modeling cognition and culture. *Journal of Cognition and Culture*, 2, 87–112.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., & Cognitive Search Research Group (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19, 46–54.
- Hume, D. (1748/2006). *An inquiry concerning human understanding*. Fairford UK: Echo Library
- Jacobmeier, D. (2004). Multidimensional consensus model on a Barabási–Albert network. *International Journal of Modern Physics C*, 16, 633–646.
- Kane, A. A., Argote, L., & Levine, J. (2005). Knowledge transfer between groups via personnel rotation: Effects of social identity and knowledge quality. *Organizational Behavior and Human Decision Processes*, 96, 56–71.
- Kummerfeld, E. & Zollman, K. J. S. (2016). Conservatism and the scientific state of nature. *British Journal for the Philosophy of Science*, 82, 956–968.
- Lorenz, J. (2003). *Mehrdimensionale Meinungsdynamik bei wechselndem Vertrauen*. Diploma thesis, University of Bremen, available at <http://nbn-resolving.de/urn:nbn:de:gbv:46-dip1000000564>.
- 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.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2, 71–87.
- Mutukrishna, M., Morgan, T. J. H., & Henrich, J. (2016). The when and who of social learning. *Evolution and Human Behavior*, 37, 10–20.
- O'Connor, C. & Weatherall, J. O. (2019). *The misinformation age: How false beliefs spread*. New Haven CT: Yale University Press
- Olsson, E. J. (2008). Knowledge, truth, and bullshit: Reflections on Frankfurt. *Midwest Studies in Philosophy*, 32, 94–110.

- Pluchino, A., Latora, V., & Rapisarda, A. (2006). Compromise and synchronization in opinion dynamics. *European Physical Journal B*, 50, 169–176.
- Richerson, P. J. (2019). An integrated bayesian theory of phenotypic flexibility. *Behavioural Processes*, 161, 54–64.
- Riegler, A. & Douven, I. (2009). Extending the Hegselmann–Krause model III: From single beliefs to complex belief states. *Episteme*, 6, 145–163.
- Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In epistemic networks, is less really more? *Philosophy of Science*, 84, 234–252.
- Schawe, H., Fontaine, S., & Hernández, L. (2021). Network bridges foster consensus: Bounded confidence models in networked societies. *Physical Review Research*, 3, 023208.
- Schurz, G. (2008). The meta-inductivist's winning strategy in the prediction game: A new approach to Hume's problem. *Philosophy of Science*, 75, 278–305.
- Schurz, G. (2019). *Hume's problem solved: The optimality of meta-induction*. Cambridge MA: MIT Press.
- Schurz, G. & Hertwig, R. (2019). Cognitive success. *Topics in Cognitive Science*, 11, 7–36.
- Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1, 43–62.
- Sunstein, C. R. (2019). *Conformity: The power of social influences*. New York: New York University Press.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge MA: Harvard University Press.
- Tomasello, M. (2019). *Becoming human: A theory of ontogeny*. Cambridge MA: Harvard University Press.
- van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford: Oxford University Press.
- Wolpert, D. H. (1996). The lack of a priori distinctions between learning algorithms. *Neural Computation*, 8, 1341–1390.
- Wolpert, D. H. & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1, 67–82.
- Wood, L. A., Kendal, R. L., & Flynn, E. G. (2012). Context-dependent model-based biases in cultural transmission: Children's imitation is affected by model age over model knowledge state. *Evolution and Human Behavior*, 33, 387–394.
- Wood, L. A., Kendal, R. L., & Flynn, E. G. (2013). Whom do children copy? Model-based biases in social learning. *Developmental Review*, 33, 341–356.
- Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74, 574–587.
- Zollman, K. J. S. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72, 17–35.
- Zollman, K. J. S. (2011). Social network structure and the achievement of consensus. *Politics, Philosophy & Economics*, 11, 26–44.