



Epistemic diversity and industrial selection bias

Manuela Fernández Pinto¹ · Daniel Fernández Pinto²

Received: 28 July 2021 / Accepted: 11 April 2023 / Published online: 15 May 2023
© The Author(s) 2023, corrected publication 2023

Abstract

Philosophers of science have argued that epistemic diversity is an asset for the production of scientific knowledge, guarding against the effects of biases, among other advantages. The growing privatization of scientific research, on the contrary, has raised important concerns for philosophers of science, especially with respect to the growing sources of biases in research that it seems to promote. Recently, Holman and Bruner (2017) have shown, using a modified version of Zollman (2010) social network model, that an industrial selection bias can emerge in a scientific community, without corrupting any individual scientist, if the community is epistemically diverse. In this paper, we examine the strength of industrial selection using a reinforcement learning model, which simulates the process of industrial decision-making when allocating funding to scientific projects. Contrary to Holman and Bruner's model, in which the probability of success of the agents when performing an action is given a priori, in our model the industry learns about the success rate of individual scientists and updates the probability of success on each round. The results of our simulations show that even without previous knowledge of the probability of success of an individual scientist, the industry is still able to disrupt scientific consensus. In fact, the more epistemically diverse the scientific community, the easier it is for the industry to move scientific consensus to the opposite conclusion. Interestingly, our model also shows that having a random funding agent seems to effectively counteract industrial selection bias. Accordingly, we consider the random allocation of funding for research projects as a strategy to counteract industrial selection bias, avoiding commercial exploitation of epistemically diverse communities.

Keywords Industrial selection bias · Epistemic diversity · Privatization of science

✉ Manuela Fernández Pinto
m.fernandezp@uniandes.edu.co

Daniel Fernández Pinto
dfernandezp@unal.edu.co

¹ Department of Philosophy and Center of Applied Ethics, Universidad de los Andes, Cra 1 No. 18a-12, G531, Bogotá, Colombia

² Independent researcher, Valencia, Spain

1 Introduction

Philosophers of science have argued that diversity brings not only political benefits to scientific communities, but also epistemic benefits. One important reason for this, the argument goes, is that having a diversity of methodological and theoretical approaches, as well as diverse background assumptions, allows the scientific community to filter different biases which might remain hidden in a homogenous group. In this sense, this epistemic diversity contributes to increasing the objectivity of the scientific knowledge produced by the community as a whole (Intemann, 2009). Furthermore, this epistemic diversity is taken to arise, at least partially, from demographic diversity, i.e., diversity in terms of the social background-race, gender, wealth, politics, culture, religion, etc.-of members of the community. In this sense, having a more demographically diverse scientific community is (frequently) expected to produce epistemically diverse approaches to problem-solving (Peters, 2021, p. 33). Longino, among other, has defended this view, claiming that a community with effective criticism, and in particular, one with tempered equality of intellectual authority, in which members of diverse social locations are included, is necessary for identifying idiosyncratic biases and producing scientific knowledge:

...the greater the number of different points of view included in a given community, the more likely it is that its scientific practice will be objective, that is, that it will result in descriptions and explanations of natural processes that are more reliable in the sense of less characterized by idiosyncratic subjective preferences of community members than would otherwise be the case. (Longino, 1990, p. 80)

Accordingly, scientific institutions should welcome members of diverse demographic backgrounds not only because the inclusion of historically marginalized individuals from science is the right thing to do morally or politically, but also because demographically diverse groups are expected to increase the epistemic diversity in the scientific community and thus produce better science.

Thus, philosophers of science have argued that epistemic diversity is an epistemic asset for science (see also, Harding, 1986; Haraway, 1989; Solomon, 2001; Rolin, 2002), guarding against the effects of biases, among other advantages. The growing privatization of scientific research, on the contrary, has raised important concerns for philosophers of science, especially with respect to the growing sources of biases in research that it seems to promote. Financial conflicts of interest in particular seem to introduce biases in the research process, leading to biased results. As the latest meta-analyses have systematically showed, industry-sponsored studies are significantly more likely to obtain results favoring sponsors than independently-funded research (Bekelman et al., 2003; Lexchin et al., 2003; Sismondo, 2008; Lundh et al., 2017). Surprisingly, the same meta-analyses have also shown that industry-sponsored studies have lower risk of bias (e.g. from blinding), and their methodological quality is at least as good as, sometimes even better than, the quality of independent studies.

This seemingly contradictory result, tells us, among other things, that the biasing mechanism of industrial funding is not related to scientific misconduct, i.e., industry sponsored research is not obtaining favorable results due to plagiarism, fabrication

or falsification of data. On the contrary, according to the latest meta-analyses, when evaluating industry sponsored research through the available controls for research quality, industry sponsored research turns out to be good quality research. We know then that financial conflicts of interest do not lead systematically to scientific fraud, but how exactly biased results enter into scientific practices is still a puzzle. Therefore, a more detailed understanding of biasing mechanisms is needed. Recently, Holman and Bruner (2017) have shown, using a modified version of a social network model developed by Zollman (2010), that industry funding can bias a scientific community without corrupting any individual scientist, especially when the community is epistemically diverse. They call this type of bias, “industrial selection bias” (Holman & Bruner, 2017, p. 1008). Their result illuminates an important mechanism through which bias can be introduced in the scientific process, despite (and also taking advantage of) the community’s epistemic diversity. In this way, Holman and Bruner argue against, or at least show a limitation of, the previously presented view that diversity contributes to expunging biases from scientific communities.

In this paper, we examine the strength of this industrial selection bias using a reinforcement learning model, which simulates the process of industrial decision-making when allocating funding to scientific projects. Contrary to Holman and Bruner’s model, where the probability of success of the agents when performing an action is given a priori, in our model the industry learns about the success rate of individual scientists and updates the probability of success on each round. The results of our simulations show that even without previous knowledge of the probability of success of an individual scientist, the industry is still able to disrupt scientific consensus, i.e., the scientific community’s convergence to the same belief. In fact, and also consistent with Holman and Bruner’s results, the more epistemically diverse the scientific community, the easier it is for the industry to move scientific consensus to the opposite conclusion. Interestingly, our model also shows that having a random funding agent seems to effectively counteract industrial selection bias. Thus, we consider the random allocation of funding for research projects as a strategy to counteract industrial selection bias, avoiding commercial exploitation of epistemically diverse communities.

The paper is divided as follows. The second section explains the problem of industrial bias that has appeared with the increasing privatization of clinical research. The third section presents Holman and Bruner’s model of industrial selection bias, in which they explore a possible mechanism behind industrial bias. In the fourth section, we analyze some limitations of Holman and Bruner’s model, which allow us to introduce our own reinforcement learning model of industrial selection in section five. In section six we explain and discuss the results of our simulations.

2 Industrial bias in clinical research

While the latest meta-analyses show that the private funding of clinical research in fact biases research results, the mechanisms through which this industrial bias operates are less clear. Biased methodological decisions that might go unnoticed through standard bias checks, include: (i) decisions regarding experimental design, such as

selection of inappropriate comparators or inappropriate doses of comparators (Djulbegovic et al., 2003), selection of poor surrogate outcomes (Bekelman et al., 2003), inadequate double-blinding (Lundh et al., 2017), and larger study sizes (Booth et al., 2008), (ii) decisions regarding the interpretation and presentation of results, such as over-interpreting results and use of spin in conclusions (Boutron et al., 2010), and (iii) decisions regarding publications, such as multiple publications of the same results, not publishing negative results (Schott et al., 2010), and ghostwriting tactics (Sismondo, 2008); all of which have focused on individual decision-making. Accordingly, the main explanation for industrial bias in clinical research has focused on different methodological decisions made by individual scientists. Less has been said regarding strategic decisions that industrial sponsors can make to shape research results in the aggregate.

Recently, Holman and Bruner (2017) have shown that an industrial selection bias can emerge in a scientific community without corrupting any individual scientist. Using a version of Zollman's social network model of scientific communication (2010), they argue that industry favorable outcomes, such as the ones detected in meta-analyses, can be produced, without corrupting any scientist in particular, if three conditions are present: (1) The community is epistemically diverse, (2) there is a merit-based structure, and (3) industry can distribute resources selectively. Accordingly, for the industry to influence scientific results and consensus, it does not need to influence directly the methodological decisions of any individual scientist, but just select and fund those scientists who already have industry-friendly views.

The phenomenon of industrial selection bias is extremely worrisome, given that it seems to extirpate any responsibility from scientific agents and relocates the moral responsibility in industrial decision making. Thus, it is important to understand the scope of and the mechanisms operating behind such phenomenon. In other words, if industrial selection bias can account for the industry bias that meta-analyses detect, even partially, this finding can have important repercussions for our understanding of privately funded research as well as for how to address industry bias.

Moreover, given that we cannot measure industry bias directly, (for example by asking pharmaceutical companies for their strategies to produce favorable research), and that the available evidence either is just unable to detect biasing mechanisms (RCTs) or is too specific to make reliable generalizations (case-by-case analyses), we consider that simulations, such as Holman and Bruner's, are an important tool here [despite their shortcomings (Martini & Fernández Pinto, 2016)], precisely because empirical research has important limitations for uncovering these mechanisms. The models also help us to explore the impact of possible countering measures before implementing them, and although they cannot predict in detail what will happen in real case scenarios, they can give us good reasons to choose the implementation of some counteracting measures over others.

3 Industrial selection bias using social networks

Based on a previous computer simulation model of the structure of communication in scientific communities based on social networks (Zollman, 2010), Holman and Bruner (2017) have developed a similar model to simulate industrial selection bias. They aim to

show that industrial selection can be the primary operating mechanism through which the industry influences the results of scientific research, without compromising the integrity of any individual scientist. In other words, scientists may not be intentionally producing industry-friendly science.

Zollman's model is based in turn on an economic model (Bala & Goyal, 1998) on bandit problems and social networks. Individual agents have to decide between two actions (A and B), each with certain probability of success (p_A and p_B). Individuals perform the action they consider more likely to be successful, and then they update their beliefs according both to the results of their action and the results of the actions of other agents in their social network. Altering the structure of the agent's social network, i.e., the number of agents she has contact with, allows to model different patterns of information sharing in a scientific community. In this way, Zollman aims to show which is the optimal way of sharing information within scientific communities (networks), and in particular how much sharing is optimal.

Holman and Bruner (2017) use Zollman's social network model to explore a different problem, i.e., the effects of industry funding on scientific research. They motivate their model through the example of industrial funding of antiarrhythmic drugs research. Given that arrhythmias frequently lead to heart attacks, it was (wrongly) thought by some researchers that suppressing arrhythmias would help prevent heart attacks, while other researchers remained skeptical of this causal hypothesis (Holman & Bruner, 2017, p. 1010). While none of these researchers conducted ethically questionable research, it was easy for the industry, as Holman and Bruner explain, to decide to fund those researchers who were more likely to develop marketable medication. When results on the dangers of antiarrhythmic drugs surfaced, it was already late for thousands of patients whose lives were lost after pharmaceutical companies made such medication available (for more details, see Holman & Bruner, 2017, pp. 1010–1012).¹

In order to simulate the industry's selection of friendly research, Holman and Bruner introduce three changes to the original model. First, in the new model not all researchers are equally productive, second, researchers have different probability of success when performing the same action, and third, the model incorporates an entry/exit dynamic, which simulates scientists entering and leaving their academic communities (2017, p. 1013).

Let us turn our attention to the second feature of the model. Given that different scientists might be using different methodologies to test a hypothesis, e.g., when testing hypothesis A , their probability of confirming A might differ when performing the same action to test A . The success rate, where success is understood as confirming A , of an

¹ Another infamous example how the industry might support favorable research without corrupting individual scientist is the tobacco industry's financial support of research on the genetic causes of cancer (Proctor, 2011). More generally, methodological biases can indeed be independent from researchers' skills and integrity. For example, research on a catalysis reaction, "a protocol for Suzuki–Miyaura type cross-coupling reactions in the presence of an iron salt," was later retracted due to other research groups' inability to replicate the findings. As in many other cases of research on catalysts, it turned out that the original results were due to the quality of the reagent used. This particular reagent had minuscule traces of metal impurities which were initially undetected by the researchers, but which turned out to be crucial for the resulting chemical reaction (Thomé et al., 2012). As this case reveals, the researchers were introducing a methodological bias in their study (stemming from their chemical supplier), which was in no way due to unethical conduct or lack of skill on their part.

individual i , when performing an experiment to test hypothesis A (p_A^i) is determined by a draw from a normal distribution centered at p with variance σ^2 . Accordingly, p_A^i accounts for a methodological bias that the individual scientist i introduces when researching.

This methodological bias is an important feature of the model, given that it is here where we can find the epistemic diversity of the community. Different individual scientists or research groups will have different ways of conducting research. This might be due to their particular background assumptions, education, methodological strengths, available equipment, and so forth. In this sense, the higher the variance on the distribution of the probability of success, the more epistemically diverse the community is.² Notice also that the probability of success reflects whether the methodological approach is truth-conducive, but scientists cannot know this a priori. They simply conduct their research as they consider best, and some of them, as it turns out, are making more successful decisions in this respect than others. In this sense, scientists might be conducting research using a mistaken methodological approach, without being corrupt.

The industry establishes an apparent efficacy threshold T , according to its desired outcomes, and decides to fund an individual scientist i , if her methodological bias (p_A^i) is above T , or in other words, if the scientist's methodology is more likely to obtain the desired industry-friendly outcomes. Agents who receive industry funding, increase their productivity by F . In this way, researchers who endorse a methodology that produces industry-friendly results receive industry funding and increase their productivity. On the contrary, researchers who do not endorse methodologies that lead to industry-friendly results, do not receive funding from the industry, and their productivity does not increase.

After running simulations with epistemic communities of 20 agents, Holman and Bruner find that industry funding quickly moves the community away from consensus towards the truth (or in Holman and Bruner's terms, consensus towards "the superior act"), e.g., the successful treatment, even with low values of F and T (see Fig. 1).

They conclude that, when there is enough methodological diversity, the industry can effectively manipulate the community away from the successful action through industrial selection. Furthermore, they also show that funding from an independent agency, such as the NSF, based on a meritocratic policy not only fails to counteract industrial selection bias, but actually makes it stronger. Given that industry funded

² Notice that Zollman (2010) interprets diversity differently. First, he calls it "cognitive diversity" instead of "epistemic diversity," although we think these terms are interchangeable in this context. Second, and more substantially, Zollman distinguishes between two ways of introducing diversity in his network model: "by limiting information or by endowing the scientists with extreme beliefs" (p. 17), and he concludes that if both of them are present, cognitive diversity remains indefinitely as a feature of the networks, so that they fail to converge to the truth. In this sense, he claims that diversity per se is not a desirable feature of scientific communities, but transient diversity is. Zollman's interesting analysis of the role of diversity in scientific communities is in fact different from ours. In particular, we follow Holman and Bruner, in whose model diversity is featured through increasing variance. Whether this understanding of diversity is interpreted in terms of background assumptions, methodologies, or other features, is something completely open. We have decided to interpret it as methodological diversity, following Holman and Bruner, whose model is our main interest, but we could have introduced another interpretation. The important thing here is that they are numerically the same thing. We thank one anonymous reviewer for highlighting the importance of this point to us.




Fig. 1 Proportion of individuals performing the superior act as a function of the variance of the normal distribution that initially sets the methodological biases of agents in the community. Simulations are of 20-person epistemic communities, $p_A = 0.5$, $p_B = 0.45$, $T = 0.03$, $F = 0$ (top), $F = 20$ (middle), $F = 100$ (bottom) (Holman & Bruner, 2017, p. 1014)

scientists are more productive, an independent agency looking at productivity to allocate funds meritocratically, will end up favoring researchers who have already been selected by the industry due to their methodological bias.³

4 The limits of the model

In Holman and Bruner’s model, the industry is given in advance the probability of success of each individual scientist when performing an action (p_A^i), and then makes funding decisions according to this probability and the industry threshold T . We consider this to be an important weakness in the initial assumptions of the model. To begin, it is impossible for any agent to acquire this kind of information in real life.⁴ Not even an individual scientist knows her probability of success when starting an experiment, and the probability of success of one experiment is independent from the success of previous experiments, so it cannot be inferred from previous information. Even if one tries to infer the success rate from previous results, such statistical inference would render variable results with certain probable error (just like loaded dice). In fact, if

³ Holman and Bruner also run simulations for an independent agent who follows a “selective policy,” according to which industry-funded projects are not taken into account when evaluating the productivity of an individual scientist. Following this selective policy does help to counter the industry’s biasing mechanism. However, Holman and Bruner do not explain how the independent agent is capable of identifying which research has industry funding in order to exclude it. In their model, this information is just given to the independent agent in advance. But in a world in which the majority of scientific research is both funded and performed in the private sector (NSB, 2020) and in which disclosures of financial conflicts of interests are seriously flawed (de Melo-Martín & Intemann, 2009), one should wonder how an independent funding agent can accurately collect such information.

⁴ Unless of course we are considering cases of fraud in which experiments are manipulated to render certain results. However, as we explained in the introduction, we are not interested here in such cases.

one admits there are previous results, then there is already a consensus forming within the scientific community. In this sense, discussions within the scientific community are already in progress at any point in which the industry decides to fund a particular scientist. However, the model does not capture this fact, it just assigns different probabilities of success to different scientists as if discussions on a research topic were at a starting point. Assuming that the industry knows the probability of success of individual scientists seems thus an ungranted idealization of the model.

Second, it seems plausible to think that the idealization is driving the industrial selection bias that results from Holman and Bruner's simulations. If the industry has a priori knowledge of the success rates of individual scientists, of course it can easily fund scientists to move the community away from the superior action, for it already knows in advance how successful each of its industry-sponsored scientists will be. However, this is not what actually happens with research funding schemes, where neither sponsors nor scientists really know how successful a particular experiment will be.

What if the industry is not given the success rate of individual scientists, but instead has to make funding decisions and learn about individual scientists' success as it goes? What would be the optimal policy for the industry to allocate funding in this case? And more importantly, would industrial selection bias emerge so easily if the industry does not have a priori knowledge of the probability of success of individual scientists? In order to answer these questions, we propose a new model of industrial selection based on reinforcement learning.

5 Industrial selection in reinforcement learning

In order to simulate funding agents that do not have a priori knowledge of the probability of success of individual scientists, but instead learn about scientists' success rate from their decision-making process, we decided to complement Zollman's social network model with a reinforcement learning (RL) model. RL models are particularly well-suited for our purposes because they help simulate goal-directed interactions of an agent with the environment. They have been successfully used to model games such as Go and Jeopardy, self-driving cars, drones, etc. The agent learns through a trial and error search, making decisions and updating their expected rewards as they go. All RL models have an exploration/exploitation trade-off, which means that they have to balance how much terrain they explore and how much they exploit. The more they explore, the less they exploit and vice versa. Sutton and Barto describe the main rationale behind RL models as follows:

...the basic idea is simply to capture the most important aspects of the real problem facing a learning agent interacting over time with its environment to achieve a goal. Clearly, such an agent must be able to sense the state of the environment to some extent and must be able to take actions that affect the state. The agent also must have a goal or goals relating to the state of the environment. The formulation is intended to include just these three aspects- sensation, action,

and goal-in their simplest possible forms without trivializing any of them. (Sutton & Barto, 2014, p. 4)

RL problems are formally described as Markov decision processes (MDPs), which help model time-discrete stochastic processes where outcomes have certain probability of success given an agent's decision. An $MDP = \langle S, A, P, R \rangle$, where S is a set of states, A is a set of actions, P is a set of transition probabilities, and R is a set of expected immediate rewards. Most solution methods for MDPs are based on Bellman's optimality equation. For the purposes of this simulation, we used a type of RL model called a Q -learning model, i.e., a model-free learning technique used to find the optimal action or selection policy using a Q -function (Heidrich-Meisnerr et al., 2007). Also, Q -learning models are set to converge, which is an optimal feature given the problem at hand. In Q -learning, P and R are unknown, $MDP = \langle s, a, r, s' \rangle$, and the solution is found through the Q -learning algorithm:

$$Q(s, a) \leftarrow Q(s, a) + \alpha \left(r + \gamma \max_{a' \in A} Q(s', a') - Q(s, a) \right)$$

The goal of the model is to construct a table with reward values for all s and a values, such that the action with the highest cumulative reward can be chosen at any given time. We use a $Q(s, a)$ to update all table values, as new information about the s and a is gathered. In this case, s has values in the range [0,100) while a can either be 0 (not fund) or 1 (fund). The table contains rewards for all combinations of scientists and actions. In other words, we want to know what is the best selection policy for a funding agent to maximize its reward, without knowing in advance the probability of success of independent scientists.⁵

Our aim is to model the behavior of funding agencies. Accordingly, we define two types of agent: an industry agent that has a certain commercial bias, i.e., favors policies that maximize their commercial profits, and a random agent with no bias, i.e., which does not favor any particular policy, and thus does not "learn" how to maximize her interests. Agents have to choose between two actions, i.e., to fund or not to fund a particular scientist. Notice that the agents in our RL model are the funding agencies and not the individual scientists, scientists could be better understood as pieces on the board of a game, who are moved by agents in their decision-making process. If scientists get funding, then they conduct research to confirm a hypothesis A , with some probability of success p_A^i . As Holman and Bruner (2017), we assume that A is true, and $\neg A$ is favorable to the industry.⁶

Now, our RL model works in tandem with a social network model. In particular, we use the social network model for the updating of scientists' beliefs according to

⁵ The code for our simulation model can be found at: https://github.com/danielfppps/epistemic_diversity_RL.

⁶ In this sense, our model has the underlying assumption that scientists can be better or worse at capturing what is true of the phenomena, and the probability of success is inclined towards what we set as true. We understand this is a big assumption in our model, and it is important to notice that while our model does try to represent more accurate conditions than Holman and Bruner's model, at least in so far as the industry does not know in advance the probability of success of each scientist, we are in no way close to representing the complexity of real scientific practice. We can only reiterate that our simulations can uncover potential mechanisms used by industrial agents to bias science in their favor.

research results after each round.⁷ In our RL model, the industry agent constructs a table where it holds a reward value for funding or not funding each scientist. The evaluation of each scientist is a “state” of the Q -learning model, a simulation with 100 scientists is in effect treated as a Q -learning problem with 100 states, where two actions are possible per state. This table is updated using the Q -learning equation as new information for that scientist’s experimental output is recorded after each round. If a scientist conducts research that favors the industry’s desired conclusion, then the industry agent becomes more likely to fund this scientist in the future. If it reaches an unfavorable conclusion for the industry, then funding for this scientist is less likely. The industry agent only updates its beliefs for scientists that it has funded. If a scientist has not received funding, then the industry agent does not learn anything from it during that round.

We give the following initial parameters to the model. We create 100 scientists, each assigned a value p_A^i drawn randomly from a normal distribution with mean 0.55 and certain standard deviation, which accounts for the methodological bias driver. We ran simulations for 400 cycles. As for the parameters of the agents, we give them an $\alpha = 1.05$ (forgetting factor), $\gamma = 0.5$ (discount factor) with an exploration probability of 0.5 (i.e., half the time the agents make a random choice). We chose these parameters for the scientists to achieve the optimal policy in a reasonable number of cycles. Although the alpha and gamma factor values do not seem to affect the model beyond the time to converge, a more robust exploration of their effect is an interesting topic for future work. Each agent is also able to fund as many scientist as they like for every cycle of the game (economic power = number of scientists).

The apparently large value of the exploration probability is necessary for efficient learning. This is because of the limits of statistical inference over small sample sizes. If a scientist appears to be favorable to industry within the first few cycles and the industry agent does not explore, it will effectively overcommit to this scientist, while failing to commit to scientists that might be better long term choices. Thus, industry agents need to allow some randomness in order to effectively exploit the scientific community. An industry agent that does not allow for randomness and starts with no a priori knowledge of scientists’ probabilities of success will fail to exploit the scientific community.

Each scientist also starts the simulation with two parameters, a and b , both drawn randomly from values between 0 and 1000. These represent their initial belief in either hypothesis. These values are updated at the end of each round for all scientists, adding 1 to a for every scientist that performed research that concluded A and adding 1 to b for every scientist that concluded $\neg A$. The belief for each and all scientists at the end of each round is calculated using the mode of a beta distribution⁸:

$$p = 100 \times (a + An - 1) / (a + b + nt - 2),$$

⁷ In this sense, the RL model can be understood as a complement or an addition to the social network model.

⁸ The mean of a beta distribution can be abnormally skewed away from the mode due to the potentially large contribution of outliers in this type of distribution. The mode gives us a view that is more focused on the view of the majority, rather than the mean of all views. We therefore believe that the mode would be more representative of what would establish the belief of a scientist.

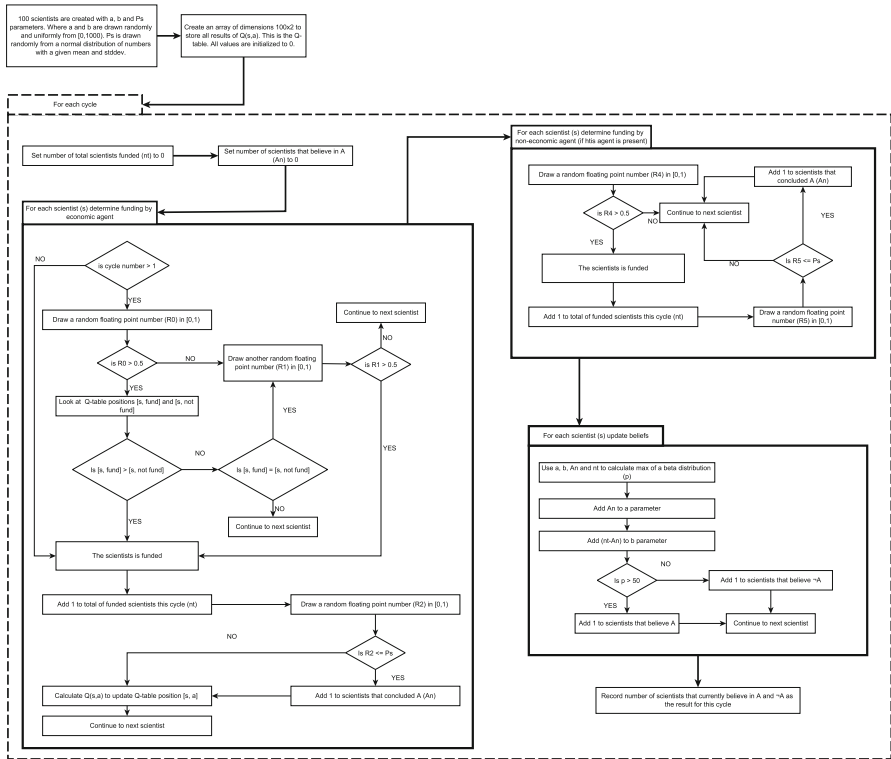


Fig. 2 Flowchart representing the coding of the simulation

where An is the number of funded scientists that favored A that round and nt is the total number of scientists that were funded in that cycle. If $p > 50$ then a scientist will believe in A while if $p \leq 50$ then the scientist will believe in $\neg A$. However, the value of p_A^i for each scientist, the methodological bias driver, remains unaffected through the entire simulation. The entire simulation process is also described in Fig. 2

6 Simulation results and discussion

We ran simulations for an industry agent, for a random agent, and for a combination of an industry agent and a random agent together. As becomes clear in Fig. 3, when the industry agent is the only funding source, scientists rapidly move away from the correct hypothesis A , and converge in the inferior, but industry-friendly outcome $\neg A$. On the other hand, when the random agent is the only source of funding, scientists are more successful at confirming the correct hypothesis, although the percentage of scientists finding the correct result decreases as they become more methodologically diverse (which is expected).

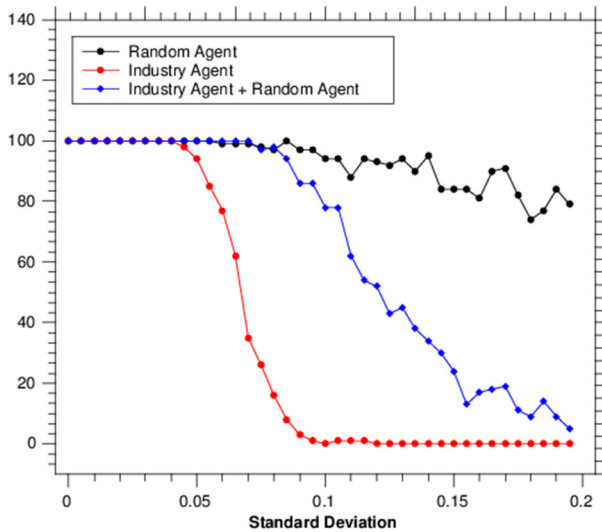


Fig. 3 Percentage of scientists preferring the correct hypothesis A with industry funding, random funding, and a combination of industry and random funding

When both industrial and random agents are funding research, we found that the random agent delays industrial selection bias, so that scientists need to be more methodologically diverse in order for the industry to pull the consensus towards the incorrect hypothesis $\neg A$. We can see industrial selection bias more clearly in Figs. 4, 5, and 6, which show the percentage of agents preferring A given only random funding, both random and industry funding, and only industry funding respectively, in 400 cycles.

In a similar vein to Holman and Bruner's results, and contrary to our initial expectations, our Q -learning model shows that even without previous knowledge of p_A^i the industry is able to effectively disrupt scientific consensus. In fact, with high variance, the industry can even move scientific consensus to the opposite conclusion, as Fig. 6 shows. This is also consistent with Holman and Bruner's earlier model, where the power of the industry to influence scientific consensus turned out to be dependent on the methodological diversity of the community: "Accordingly, as methodological diversity increases, the community is less likely to converge on the superior act" (Holman & Bruner, 2017, p. 1015). This is an important and somewhat unexpected result, given that methodological diversity has been strongly supported as a mechanism for countering bias in science (see e.g., Longino, 2002; Solomon, 2001). Without rejecting the arguments in favor of increasing diversity in scientific research, we consider that both models of industrial selection bias suggest that diversity can be co-opted in favor of commercial interests and this is a problem that should not be overlooked. In other words, we are not arguing that diversity isn't good for scientific communities, but rather that diversity can be co-opted as a strategy to favor partisan interests (see also, Fernández Pinto, 2018).

In their social network model, Holman and Bruner also examine the possibility of countering industrial selection bias through the influence of independent sources of

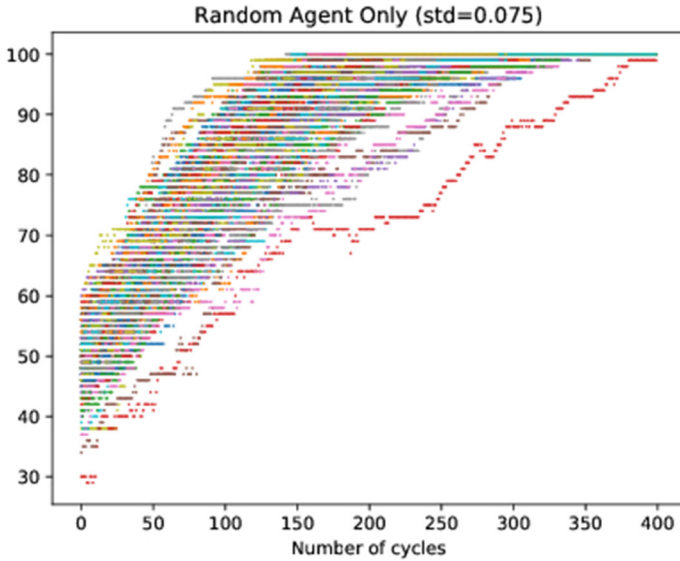


Fig. 4 Percentage of scientists preferring A, when funded exclusively by a random agent. Each curve in the image represents a 400 cycle simulation. In all simulations, all scientists believe in hypothesis A at the end of the simulation

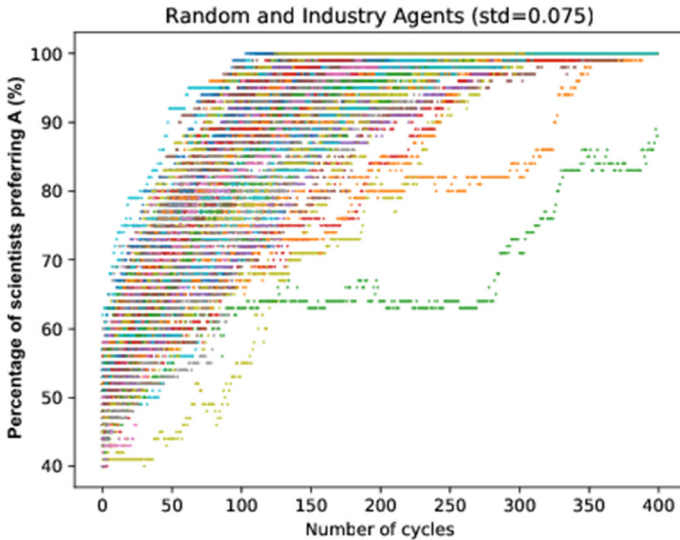


Fig. 5 Percentage of scientists preferring A, when funded by random and industry agents. Each curve in the image represents a 400 cycle simulation. It takes longer on average for scientists to reach full consensus on A, and in some simulations this does not happen at all

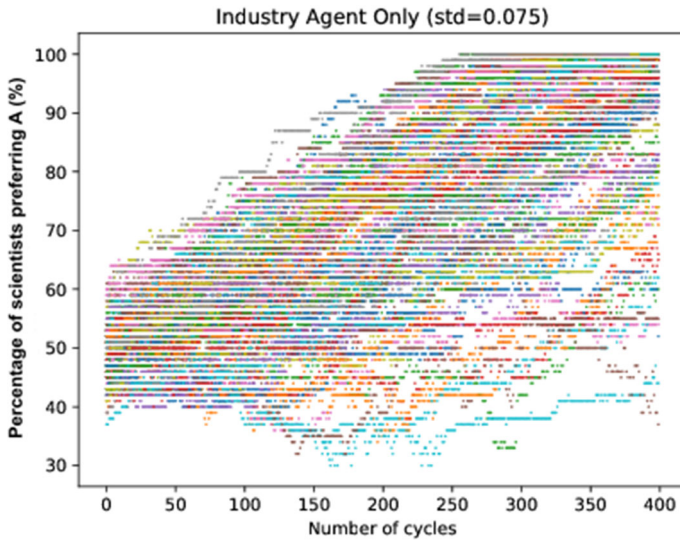


Fig. 6 Percentage of scientists preferring A, when funded exclusively by an industry agent. Each curve in the image represents a 400 cycle simulation. In many cases, the industry agent is able to completely prevent consensus on A

funding, such as the National Science Foundation, which allocate funding on a merit-based system. The initial rationale is that independent sources will favor scientific research on a meritocratic system and not on future commercial gain, helping to counteract industry bias. Holman and Bruner found, however, that independent funding might actually compromise the community further, given that independent agencies would end up “disproportionately funding individuals with industry-favorable biases” (2017, p. 1017). The reason is that industry-friendly researchers would become in general more productive because they have more research resources, obtaining higher evaluations on a merit-based system. So, in general, independent funding agencies working with a meritocratic system will not help counteract industrial selection bias.

In our simulation we did not contrast industry funding with a merit-based system, but instead we implemented an agent who funds research projects randomly. As Fig. 3 shows, having a random funding agent seems to effectively counteract industrial selection. In particular, having a random allocation of funding seems to obstruct industrial selection, so that more methodological diversity is needed for industry funding to move consensus towards the opposite conclusion. Our results suggest then that a random allocation of funding, even a partial one, might be a much better strategy to counteract industrial selection bias than using independent funding agencies based on a meritocratic system. Although this suggestion might be controversial to some, it is coherent with recent research results in social epistemology and the social economy of science, showing the advantages of a random allocation of research funding, in particular as a way of counteracting bias and discrimination in science (see e.g., Avin, 2018; Fang & Casadevall, 2016; Gross & Bergstrom, 2019; Roumbanis, 2019). We consider our simulation results one more piece of evidence in favor of such policies.

In sum, the privatization of scientific research has introduced new challenges to the epistemic goals of science. In particular, Holman and Bruner (2017) have identified industrial selection as a mechanism through which the private industry can influence scientific results without corrupting any individual scientists. They have also shown that this mechanism is more successful, the more methodologically diverse the scientific community is. In this sense, Holman and Bruner's argument shows an important limitation of the assumed benefits of methodological diversity in scientific communities that many philosophers of science have emphasized (see e.g., Haraway, 1989; Harding, 1986; Longino, 1990; Rolin, 2002; Solomon, 2001). Given industrial selection, the benefits of diversity need to be qualified. In particular, the extent to which diversity can be used to the advantage of private and commercial interests needs to be better understood in order to counteract such a manipulation and misuse of diversity efforts in science. In addition, Holman and Bruner (2017) also show that meritocratic systems of resource allocation do not seem to counteract industrial selection, but instead they even contribute to its success. Although we do not explore a meritocratic system, our results show that perhaps a system of random resource allocation would be worth exploring as a mechanism for countering industrial selection in science.

Before the concluding remarks, let us say a little more about the random allocation of funding and what it entails. Of course, many have defended a meritocratic system of resource allocation as the best way of counteracting biases and discrimination in scientific research. Holman and Bruner's results together with the results presented in this paper show that this might not be the best way to proceed. The biases that the private industry has been introducing in the development of scientific research in past decades are much more subtle and difficult to identify than what a meritocratic evaluations allows, as recent meta-analysis have clearly shown (see e.g., Lundh et al., 2017). Accordingly, the random allocation of funding emerges as a strategy worth exploring to help to counteract such biases. As Neurath reminds us, understanding the limits of rational insight is key in the process of decision-making, for those who do not recognize the limits of their own thinking and pretend they can always appeal to reason for action, may fall into a sort of pseudorationalism:

Most of our contemporaries rely on their insight and want to leave the decision in all things to it once and for all. Their starting-point is the view that given enough thought one could at least determine which manner of action has the greater probability of being successful, should certainty be impossible (...) The pseudorationalists do true rationalism a disservice if they pretend to have adequate insight exactly where strict rationalism excludes it on purely logical grounds. (Neurath, 1913, pp. 7–8)

Precisely in cases in which rational insight cannot render the desired results, the drawing of lots becomes the best available option: "If a man is no longer able to decide on the basis of insight which of several actions to prefer, he can draw lots..." (Neurath, 1913, p. 4). In a similar way, given the limits that we find in the rational allocation of research funding due to industrial selection, perhaps we should consider the drawing of lots, i.e., the random allocation of research funding, as a better option.

It is not our purpose here to defend a particular scheme for random allocation for which there are already a number of initiatives in place (see e.g., Avin, 2018; Fang &

Casadevall, 2016; Gross & Bergstrom, 2019; Roumbanis, 2019), but to suggest, following our simulation results, that it could be worth exploring as a funding strategy. It is important to clarify, however, that most strategies to allocate research funding randomly, include a two-stage process in which only the proposals that pass an initial quality filter are then added to a lottery during a second stage. In this way, the lottery system does not mean that any research proposal can get funding, but that the winners are picked from a group whose good quality and feasibility has already been established. As expected, strategies to fund research randomly have also faced important critiques (see e.g., Bedessem, 2020).

A couple of caveats. First, notice that our suggestion of exploring random allocation as a funding strategy for research takes as given the private and commercial framework of scientific organization today. If commercial interests were not a part of scientific research, perhaps a meritocratic system would work better. Second, we must emphasize that while our modeling results point to a random allocation of funding as a promising, or at least worth-exploring, alternative, further modeling, especially one that simulates the problematic factors of random allocation is needed as well, before making any commitments to this funding strategy.

7 Conclusions

The main purpose of this paper was to examine the strength of industrial selection using a reinforcement learning model. While industrial selection has already been examined by Holman and Bruner (2017), we consider their model to have an important limitation in so far as the industry is given a priori the probability of success when performing an action for each scientist, which constitutes much more information than what is reasonable to expect in real life. Accordingly, in our model the industry learns about the success rate of individual scientists and updates the probability of success on each round. The results of our simulations show that even without previous knowledge of the probability of success of an individual scientist, the industry is still able to disrupt scientific consensus. In this sense, our model corroborates what Holman and Bruner (2017) found in their simulation results. In a similar vein, we found that the more epistemically diverse the scientific community, the easier it is for the industry to move scientific consensus to the opposite conclusion. Interestingly, our model also shows that having an agent who allocates funds randomly seems to effectively counteract industrial selection bias. Thus, we suggest that the random allocation of funding for research projects would be worth examining as a strategy to counteract industrial selection bias, avoiding commercial exploitation of epistemically diverse communities.

Finally, further exploration of the parameters of our Q -learning model is still needed. As with any simulation model, we had to decide to stop the exploration at some point and collect our findings, but we acknowledge that much more is yet to be done.

Acknowledgements The authors would like to thank the editors of the special issue “Social and Cognitive Diversity in Science,” Inkeri Koskinen, Jaakko Kuorikoski, Samuli Reijula, and Kristina Rolin, for their

work and input during the review process. We would also like to thank the organizers and audiences in the Formal Social Epistemology Workshop at the University of California, Irvine, especially Cailin O'Connor, the Philosophie und Wissenschaftsreflexion Colloquium in the Institut für Philosophie at Leibniz Universität Hannover, especially Torsten Wilholt, the Biases in Science Conference at LMU, and the V Conference of the Latin American Association of Analytic Philosophy (ALFAn) where we presented preliminary versions of this article. Special thanks to Bennett Holman and Justin Bruner for sharing the code of their model with us. Finally, we thank three anonymous reviewers for their careful reading of the paper and their thoughtful comments.

Funding Open Access funding provided by Colombia Consortium.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Avin, S. (2018). Policy considerations for random allocation of research funds. *A Journal on Research Policy & Evaluation*, 1, 1–27.
- Bala, V., & Goyal, S. (1998). Learning from neighbours. *Review of Economic Studies*, 65, 565–621.
- Bedessem, B. (2020). Should we fund research randomly? An epistemological criticism of the lottery model as an alternative to peer review for the funding of science. *Research Evaluation*, 29, 150–157.
- Bekelman, J., Li, Y., & Gross, C. (2003). Scope and impact of financial conflicts of interest in biomedical research: A systematic review. *JAMA*, 289, 454–465.
- Booth, C., Cescon, D., Wang, L., Tannock, I., & Krzyzanowska, M. (2008). Evolution of the randomized controlled trial in oncology over three decades. *Journal of Clinical Oncology*, 26, 5458–5464.
- Boutron, I., Dutton, S., Ravaud, P., & Altman, D. (2010). Reporting and Interpretation of randomized controlled trials with statistically nonsignificant results for primary outcomes. *JAMA*, 303, 2058–2064.
- de Melo-Martín, I., & Intemann, K. (2009). How do disclosure policies fail? Let us count the ways. *FASEB Journal*, 23, 1638–1642.
- Djulbegovic, B., Cantor, A., & Clarke, M. (2003). The importance of the preservation of the ethical principle of equipoise in the design of clinical trials: Relative impact of the methodological quality domains on the treatment effect in randomized controlled trials. *Accountability in Research*, 10, 301–315.
- Fang, F., & Casadevall, A. (2016). Research funding: The case for a modified lottery. *MBio*, 7, 1–8.
- Fernández Pinto, M. (2018). Democratizing strategies for industry-funded medical research: A cautionary tale. *Philosophy of Science*, 85, 882–894.
- Gross, K., & Bergstrom, C. (2019). Contest models highlight inherent inefficiencies of scientific funding competitions. *PLoS Biology*, 17, e3000065.
- Haraway, D. (1989). *Primate visions: Gender, race and nature in the world of modern science*. Routledge.
- Harding, S. (1986). *The science question in feminism*. Cornell University Press.
- Heidrich-Meisner, V., Lauer, M., Igel, C., & Riedmiller, M. (2007). reinforcement learning in a Nutshell. In: *ESANN'2007 proceedings—European symposium on artificial neural networks*, 25–27 April 2007. <https://www.elen.ucl.ac.be/Proceedings/esann/esannpdf/es2007-4.pdf>
- Holman, B., & Bruner, J. (2017). Experimentation by industrial selection. *Philosophy of Science*, 84, 1008–1019.

- Intemann, K. (2009). Why diversity matters: Understanding and applying the diversity component of the National Science Foundation's broader impacts criterion. *Social Epistemology*, 23, 249–266.
- Lexchin, J., Bero, L., Djulbegovic, B., & Clark, O. (2003). Pharmaceutical industry sponsorship and research outcome and quality: Systematic review. *British Medical Journal*, 326, 1167–1170.
- Longino, H. (1990). *Science as social knowledge*. Princeton University Press.
- Longino, H. (2002). *The fate of knowledge*. Princeton University Press.
- Lundh, A., Lexchin, J., Mintzes, B., Schroll, J. B., & Bero, L. (2017). Industry sponsorship and research outcome. *Cochrane Database of Systematic Reviews*, 2, MR000033.
- Martini, C., & Fernández Pinto, M. (2016). Modeling the social organization of science: Chasing complexity through simulations. *European Journal for Philosophy of Science*, 7, 221–238.
- National Science Board. (2020). *Science and engineer indicators 2020*. National Science Foundation.
- Neurath, O. (1913). The lost wanderers of Descartes and the auxiliary motive. In R. Cohen, M. Neurath, & C. Fwacett (Eds.), *Philosophical paper, 1913–1946* (pp. 1–12). D. Riedel Pub. Co.
- Peters, U. (2021). Hidden figures: Epistemic costs and benefits of detecting (invisible) diversity in science. *European Journal for Philosophy of Science*, 11, Article Number 33.
- Proctor, R. (2011). *Golden Holocaust: Origins of the cigarette catastrophe and the case for abolition*. University of California Press.
- Rolin, K. (2002). Why gender is a relevant factor in the social epistemology of scientific inquiry. *Philosophy of Science*, 71, 880–891.
- Roumbanis, L. (2019). Peer Review or Lottery? A Critical Analysis of Two Different Forms of Decision-making Mechanisms for Allocation of Research Grants. *Science, Technology, & Human Values*, 44, 994–1019.
- Schott, G., Pacht, H., Limbach, U., Gundert-Remy, U., Ludwig, W., & Lieb, K. (2010). The Financing of Drug Trials by Pharmaceutical Companies and its Consequences. *Deutsches Aerzteblatt Online*, 107, 279–285.
- Sismondo, S. (2008). Pharmaceutical company funding and its consequences: A qualitative systematic review. *Contemporary Clinical Trials*, 29, 109–113.
- Solomon, M. (2001). *Social empiricism*. MIT.
- Sutton, A., & Barto, R. (2014). *Reinforcement learning: An introduction*. MIT.
- Thomé, I., Nijs, A., & Bolm, C. (2012). Trace metal impurities in catalysis. *Chemical Society Reviews*, 41, 979–987.
- Zollman, K. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, 72, 17–35.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.