

What is the Epistemic Function of Highly Idealized Agent-Based Models of Scientific Inquiry?

Preprint version

Daniel Frey and Dunja Šešelja¹

Abstract: In this paper we examine the epistemic value of highly idealized agent-based models (ABMs) of social aspects of scientific inquiry. On the one hand, we argue that taking the results of such simulations as informative of actual scientific inquiry is unwarranted, at least for the class of models proposed in recent literature. Moreover, we argue that a weaker approach, which takes these models as providing only ‘how-possibly’ explanations does not help to improve their epistemic value. On the other hand, we suggest that if ABMs of science underwent two types of robustness analysis, they could indeed have a clear epistemic function, namely by providing evidence for philosophical and historical hypotheses. In this sense, ABMs can obtain evidential and explanatory properties and thus be a useful tool for integrated history and philosophy of science. We illustrate our point with an example of a model—building on Zollman’s (2010) ABM—which we apply to a concrete historical case study.

Keywords: agent-based models, robustness analysis, how-possibly explanations, scientific interaction, integrated history and philosophy of science.

1 Introduction

Throughout the last decade agent-based modeling (ABM) has become one of the popular methods in formal (social) epistemology of science. The very possibility of simulating scientific inquiries has shown promise for tackling a variety of questions: from the impact of different social networks on the efficiency of knowledge acquisition (Zollman 2007, 2010, 2013; Kummerfeld and Zollman 2016; Grim 2009; Grim et al. 2013; Borg et al. 2018, 2017b) to the problem of the division of cognitive labor (Weisberg and Muldoon 2009; De Langhe 2014; Alexander, Himmelreich, and Thompson 2015; Thoma 2015; Pöyhönen 2017), to the problem of norms that are to guide peers facing a disagreement (Douven 2009, 2010; De Langhe 2013), etc. A typical feature of these models is that they are highly idealized, which allows them to offer simple explanations of complex social phenomena.² Due to the latter feature, the results of the simulations have been adopted by both philosophers of science and epistemologists (e.g. (Wray 2011; Goldman and Blanchard 2016)).

¹ The order of authors is alphabetical; both authors contributed equally to this paper. To contact the authors, please write to dunja.seselja@rub.de.

² See e.g. (Reutlinger, Hangleiter, and Hartmann 2016) for a discussion on the explanatory features of highly idealized or toy-models.

Nevertheless, recent discussions in this domain have suggested a dose of caution when adopting results obtained by means of such models. On the one hand, it has been shown that some of these findings are not robust under the changes in the relevant parameter space (Rosenstock, O'Connor, and Bruner 2017), nor are they replicated once we change some of the underlying assumptions of the models (Borg et al. 2017a, 2018, 2017b; Frey and Šešelja 2018). In addition, it has been argued that the models lack empirical calibration, which impedes the reliability of their results with respect to normative conclusions we may draw from them (Martini and Pinto 2016).

In response to these objections it has been suggested that ABMs of science should be understood as providing 'how-possibly' (rather than 'how-actually') explanations (e.g. Rosenstock, O'Connor, and Bruner 2017). In other words, rather than showing how something actually happens, they show how and under which circumstances the given phenomenon is possible (Reutlinger, Hangleiter, and Hartmann 2016; Rosenstock, O'Connor, and Bruner 2017). However, even if we grant these models the how-possibly explanatory function, it remains unclear what is the exact epistemic value of showing a possibility, and whether these models can perhaps be used for additional epistemic gains.

The aim of this paper is to address these questions. On the one hand, we will argue that the epistemic function of recently developed highly idealized ABMs of scientific interaction is only heuristic or exploratory (Reutlinger, Hangleiter, and Hartmann 2016; Gelfert 2016), rather than explanatory.³ Importantly, we will show that how-possibly explanatory function isn't epistemically valuable in case of these ABMs. On the other hand, we will suggest that this class of models can in fact become useful as a method for providing evidence for historical and philosophical hypotheses concerning scientific inquiry, and that as such they can play evidential and explanatory roles.

Here is how we will proceed. In Section 2 we will give a brief overview of Zollman's (2010) ABM as a paradigmatic example of highly idealized ABMs of scientific interaction, and show how it has been used to suggest normative claims concerning actual scientific inquiry. In Section 3 we will raise objections against such applications of these models and argue for their lack of explanatory properties. We will then turn to our proposal for how these models can be improved to become epistemically valuable. In Section 4 we suggest how ABMs can play evidential and explanatory roles in the domain of integrated history and philosophy of science. In Section 5 we illustrate our point by showing how an adjusted version of Zollman's (2010) ABM, first presented by Frey and Šešelja (2018), can be used to analyze a historical case study. As an example of such a case study we take the research on peptic ulcer disease from the history of medicine, which is examined by Zollman as well. In Section 6 we make some concluding remarks.

³ Henceforth, when speaking of the *epistemic* function of models, we will refer to a function that goes beyond their exploratory or heuristic value.

2 Zollman's (2010) ABM and its Impact

2.1 The model

Zollman motivates the design of his ABM with an example from the history of medicine. During the first half of the twentieth century the research on peptic ulcer disease (PUD) proceeded along two main lines: one based on the 'bacterial hypothesis', according to which the disease is caused by bacteria, and another one based on the 'acidity hypothesis', according to which the disease is caused by excessive acidity in the stomach. In the 1950s, after the publication of a large case study, which showed that there are no bacteria in the human stomach,⁴ the former of the two hypotheses was dropped. However, 30 years later, with Warren and Marshall's discovery of *Helicobacter pylori*, the bacterial hypothesis was revived and this bacterium came to be known as the major cause of PUD. Cases such as this one raise the question, which social-epistemic factors are likely to increase the chance of a premature abandonment of fruitful hypotheses.

Zollman examines a hypothesis that the culprit can be found in a too wide information sharing: Palmer's erroneous results were shared too widely across the medical community at the time, leading to the premature abandonment of the bacterial theory of PUD. In order to test this hypothesis, he uses an ABM inspired by Bala and Goyal (1998) and based on the idea that the process of scientific inquiry can be tackled as a type of *bandit problem*. The so-called bandit problems, usually discussed in the context of economics and game theory, concern the following question: if a gambler is confronted with different slot machines, at which point should she stop testing which machine maximizes her reward, and stick with the one that seems the best in this respect? Zollman suggests that an analogous situation occurs in the context of scientific inquiry, where we can ask: at which point should a scientist stop testing different hypotheses and stick with the most promising one? The payoff of a slot machine here corresponds to the success of applying the given hypothesis (or a method, or a theory), while the objective probability of success (OPS) of a slot machine corresponds to the OPS of the given hypothesis (or a method, or a theory). Hence, just like a gambler faces the choice between different slot machines, so does a scientist face the choice between different hypotheses.

Adding a social dimension to this problem raises the following question: if more than one gambler is trying to determine which bandit is the most profitable one, how does the information flow among the gamblers influence their respective choices? Analogously, we can ask: if more than one scientist is working in the same domain consisting of rivaling theories, how does the information flow among them influence their respective decisions as for which theory to pursue?

The model is designed as a computer simulation, which is round based.

⁴The study, conducted by the famous gastroenterologist Palmer, was based on a methodological error (see (Šešelja and Straßer 2014)).

At the start of a run each agent is assigned a random prior value for each of the two rivaling theories. Every round an agent makes 1,000 pulls, each of which can be a success or failure, where the probability of success is given by the OPS of the respective theory. Agents then update their beliefs via Bayesian reasoning (modeled by means of beta distributions), in view of their own success and the success of some other agents, namely those with whom they are connected in a social network. Zollman investigates the efficiency of agents in converging onto the hypothesis with the better OPS in three kinds of social networks: a cycle, a wheel and a complete graph (see Figure [1]⁵). In the cycle, every agent exchanges information with two of her neighbors. The wheel has the same information flow as the cycle, except that in addition one agent exchanges information with all the other agents. Finally, in the complete graph every agent exchanges information with every other agent.




Figure 1: A cycle, a wheel and a complete graph. The nodes in each graph represent agents, while the edges that connect the nodes represent transmission of information between two agents.

Assuming that the OPS of the better theory is 0.5 and the OPS of the worse theory 0.499, Zollman's results show the cycle scores the best and the complete graph the worst. More precisely, agents connected in the cycle converge on the better theory the highest number of runs, those connected in the wheel converge on the better theory a smaller number runs, and those connected in the complete graph the smallest number of runs. This suggests that information flow via highly connected groups can be epistemically harmful:

It would appear here that the amount of information distributed is negatively impacting the ability of a social group to converge on the correct methodology. Initially suggestive information is causing everyone to adopt one particular methodology. (p. 28)

However, once agents are modeled as biased towards hypotheses with which they initially start, the results become inverse.⁶ In other words, if agents are modeled as

⁵ We are grateful to AnneMarie Borg for this illustration.

⁶ Bias of agents is represented by assigning them more extreme distributions (i.e. drawing α/β for the beta distribution from a larger interval) as priors from which they start their updates, thus allowing for a more conservative outcome of the

more resistant to changing their beliefs in view of new information, misleading initial results can't infect the entire community.

Thus, Zollman's model suggests two ways of reducing errors in the process of learning of a scientific community: either the information flow needs to be restricted, or scientists need to be initially biased towards their pursued hypotheses. In the remainder of the paper we will focus on the former suggestion.

2.2 The reception and criticism

Zollman's results have been influential. For example, (Strevens 2010) calls upon Zollman's results when suggesting:

It might ... be preferable for scientists not to take into account too much information about their colleagues' beliefs about a problem, if a few authoritative pronouncements would stifle much-needed diversity in the range of approaches to a problem. ... The question how to tune attention to authority in the short term so as to find a level of diversity that maximizes correctness in the long term has been explored with considerable insight by Zollman... (p. 306).

Similarly, (Wray 2011) brings up Zollman's results when suggesting that: "Full communication within the community of scientists would be detrimental, for effective lines of communication not only serve to spread the truth; they also serve to spread falsehoods." (p. 134) Or take an entry in the *Stanford Encyclopedia of Philosophy* on Social Epistemology (Goldman and Blanchard 2016), which mentions Zollman's results, concluding: "This is another illustration of the idea that *prima facie* detrimental features of the practice of science (reduced communication and dogmatism) may in fact be epistemically beneficial."

Finally, as we have mentioned above Zollman (2010) has built the model to test an explanatory hypothesis concerning the case of research on PUD. In view of the results of his simulations he writes:

For PUD, I have suggested that things might have been better had Palmer's result not been communicated so widely or had people been sufficiently extreme in their beliefs that many remained unconvinced by his study. (p. 33)

Recently, however, a set of Zollman's models have received critical response. First, Rosenstock, O'Connor, and Bruner (2017) have shown that the so-called "Zollman effect", namely the positive difference between successful convergence in the cycle

updating process. It's worth mentioning though that the reversal of results is only due to the cutoff point at which Zollman chooses to end his simulation. If the model is run sufficiently long, agents in all networks end up on the correct theory: since extreme priors allow for the theoretical diversity to be preserved, agents in all networks have enough time to gather evidence and make the right choice.

versus the complete graph, significantly decreases or even disappears once we extend the parameter space used in the simulations. The authors suggest that the results seem to be robust only under conditions characteristic of difficult inquiry. Moreover, Borg et al. (2018, 2017b) have failed to reproduce Zollman effect in an argumentation-based ABM of scientific interaction. Finally, Frey and Šešelja (2018) show that once we de-idealize some of the assumptions in Zollman's model, Zollman effect doesn't appear even in some typical situations of *difficult* inquiry.

In view of this, applying Zollman's results to actual scientific practice seems unwarranted unless the context in the case of actual inquiry can be linked to the conditions represented in the model under which Zollman effect robustly holds. But what about the model providing at least how-possibly explanations? We now turn to this question.

3 Possibility is not enough

3.1 How-possibly explanation?

In view of the above mentioned criticism of Zollman's model it has been suggested that we should understand this ABM as providing a how-possibly explanation (Rosenstock, O'Connor, and Bruner 2017). The model, so the argument goes, doesn't show that the cycle network is superior to the complete graph in actual scientific inquiry, but only that this is a possibility. Such a possibility can be interesting in showing, for instance, that a certain prevalent opinion concerning the target phenomenon may not necessarily hold, in which case we may learn something new from the model (see (Grüne-Yanoff 2009)). For example, in case of Zollman's model, it could be argued that it casts doubt on the view that interaction among scientists is always epistemically beneficial. Similarly, we could say that it shows a possible cause of the premature abandonment of a fruitful hypothesis in the case of PUD research.

The notion of how-possibly explanations, first introduced by Dray (1957) with regard to explanations in history, has been extensively discussed in the literature on scientific modeling, especially in biology and social sciences (see e.g. Forber 2010, 2012; Reydon 2012; Bokulich 2014; Ylikoski and Aydinonat 2014; Riel 2015). Nevertheless, the question of the epistemic import of such explanations—whether they genuinely explain, and if so, what kind of explananda, or whether they should rather be considered as heuristic tools—has remained disputed. Our intention here is not to settle this debate, but rather to suggest that we should err on the side of caution when using the how-possibly strategy to defend highly idealized models of science.

The following caricatured example of a model brought up by Pfliegerer (2014) nicely illustrates the main problem with the idea that models are useful in virtue of explaining how something is possible:

To establish that high intake of alcohol is the natural (distortion free) result of human liquid-drink consumption, the model rules out liver disease, DUIs, health benefits, spousal abuse, job loss and all other distortionary factors. By positing these idealized conditions, the model obviously ignores some important determinants of human alcohol consumption in the real world. However, ...it allows a meaningful role for humans as producers of that pleasant “buzz” one gets by consuming alcohol, and shows clearly that ...it is optimal for humans to be drinking all of their waking hours.⁷

While it may be entertaining to consider a world in which it is optimal for humans to be drinking all the time, it doesn't tell us much about the effects of actual alcohol consumption. This is because the model could represent a scenario that belongs only to a mathematical or logical possibility, which is in no relevant way related to the real-world phenomena and which doesn't increase our understanding of these phenomena. This doesn't mean that in order to obtain epistemically valuable results, the model needs to represent only the real-world phenomena. To the contrary, modeling counterfactual scenarios can indeed be informative of the significance of certain factors in actual causal dependencies. If we, for example, represent a decision-making scenario in which all agents are biased, this might be informative of the role of biases in the given context.

In order to clarify the relevance of explaining counterfactual scenarios, Ylikoski and Aydinonat's (2014) discussion of understanding via how-possibly explanations will be of help. The authors distinguish between how-possibly explanations aimed at a particular empirical fact, representing a *possible causal scenario*, and how-possibly explanations addressing processes that can bring about certain kind of effects, representing *causal mechanism schemes*. The latter “do not directly explain any particular empirical fact. They address only simplified *theoretical explananda*” (p. 27, italics added). Therefore, the epistemic import of causal mechanism schemes is only indirect: while they do not explain any concrete facts, they are informative of the menu of possible explanations for a given event and of possible causal scenarios that may underlie it. As such, they may incorporate a number of counterfactual conditions. Nevertheless, both types of explanations can facilitate our understanding of the given target phenomena only if we are able to make correct inferences about them, where the key criterion of understanding is one's ability to make inferences about counterfactual situations concerning these phenomena (Ylikoski and Aydinonat 2014, p. 29).

Thus, the problem with models such as Pfliegerer's caricature is not that it represents 'just a possibility', but rather that the represented possibility provides us with no ability to make *relevant* (or interesting) counterfactual inferences. This can easily happen if a model includes too many counterfactual factors at the same time

⁷ The quote is taken from the English version of Pfliegerer's article: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.432.4076&rep=rep1&type=pdf>, p. 9, accessed on February 14, 2018.

(i.e. if it is highly idealized). On the one hand, the represented possibility may end up being entirely uninteresting (Pfleiderer's example is a case in point). On the other hand, we may have no way of making (interesting or relevant) inferences about individual causal dependencies unless there is a way to discern the impact of certain individual factors on the given result. We will now turn to the method of robustness analysis, which has been introduced precisely to tackle this kind of problems.

3.2 The importance of robustness analysis

As it is well-known in the literature on scientific modeling (see e.g. Soler et al. 2012; Lehtinen 2017), a common way to address the above problems is by confirming the robustness of results obtained by a model, under two types of changes.

First, we can speak of the robustness of results under parameter changes:⁸ by specifying the scope of parameters within which the results remain stable we can determine their significance for the given phenomenon. For instance, if the results hold only for a small portion of the parameter space, (as it has been shown for Zollman's model), we can conclude that they are likely to occur only under very limited circumstances. Which circumstances these are is determined by mapping the given parameter space to different scenarios in which the target phenomenon occurs. To take a simple example, if the results of an ABM change when the number of agents increases, then we can relate these different outcomes to different community sizes. Or as we have seen above (Section 2.2), Zollman effect has been restricted to the context of difficult inquiry by mapping the parameter space in which it holds to the specific context of inquiry. The most reliable way to conduct such a mapping is by empirically calibrating the model, where specific parameters are given by means of empirical knowledge relevant for the given phenomenon (Martini and Pinto 2016).

The second type of robustness analysis concerns the stability of results under the changes in the idealizing assumptions of the model. This is especially important in case of highly idealized models where results may easily depend on idealizations rather than on assumptions that adequately represent the target phenomenon.⁹ More precisely, if a model consisting of plausible conditions about the given target phenomenon and implausible idealizations produces phenomenon X , and we observe that when some of these implausible idealizations are replaced with more realistic assumptions concerning some context C , we fail to produce X , then the model is not informative of X in C . In such a case, X is produced in virtue of idealizing assumptions (which do not adequately represent context C). For example,

⁸ This process is usually called sensitivity analysis, see e.g. Thiele, Kurth, and Grimm 2014.

⁹ Such analysis is also known as *derivational robustness analysis*, see Kuorikoski and Lehtinen 2009; Ylikoski and Aydinonat 2014. The importance of de-idealization for the robustness analysis has been extensively discussed in the literature on models in economics, see e.g. Mäki 2009; Lehtinen and Kuorikoski 2007.

(Frey and Šešelja 2018) analyze Zollman's model, arguing that the Zollman effect occurs only under the assumption that it is neither the case that scientists reason cautiously, nor is it the case that they are in a situation where similarly successful theories count as equally good.

To sum up, the two types of robustness analysis are essential for relating models to their target phenomena since they allow us to specify the context which the model adequately represents. Note however that robustness analysis comes in degrees and hence, it may not control for all idealizations of the model. For instance, we may identify a robust property that holds across a range of parameters and changes in idealizing assumptions, which may thus increase our confidence that the same property holds for a certain real-world phenomenon as well. Nevertheless, we may not be able to test the robustness of the model with respect to all the idealizing assumptions.¹⁰ Clearly, the more thorough checks we make, the more reliable our conclusions will be. With this in mind, we turn now to the question of the epistemic function of ABMs.

4 Evidential and explanatory roles of ABMs

In this section we will examine the epistemic function of ABMs of science by focusing on the role they could play in the domain of integrated history and philosophy of science (HPS).

Let's start by recalling Zollman's original research question, concerning the historical case study on PUD: which social-epistemic factors are likely to increase the chance of a premature abandonment of fruitful hypotheses? Zollman used his ABM to test whether too wide sharing of Palmer's results led the whole community to abandon the bacterial hypothesis of PUD. While the results confirming this conjecture may not be reliable (as discussed in Section 2.2), Zollman's attempt is significant from a methodological point of view: it aims at introducing simulations as the method for obtaining evidence for hypotheses in the domain of HPS. We will now take a look at the traditional HPS method of obtaining evidence, and in view of this specify how ABMs could be of help.

4.1 The method of historical case studies

The literature on HPS has traditionally used the method of historical case studies as the primary source of evidence. Nevertheless, this approach is not without problems. Case studies may be cherry-picked to fit one's philosophical assumptions (see e.g. Pitt 2001), or they may be theory-laden in the sense that there is no unique objective narrative that underlies the reconstruction and the identification of

¹⁰ Nor do robustness checks necessarily lead to a fully de-idealized model: the given model may still employ a number of (original) idealizations, though in view of robustness analysis we may have shown that these idealizations are not problematic.

historical events (Kinzel 2015). And while this poses a limit to the overall reliability of the evidence that is obtained by means of historical case studies, it doesn't impede their evidential function. Kinzel distinguishes four different epistemic functions that historical case studies may play:

1. Novelty: "case studies can provide us with new, previously unknown and perhaps surprising information" (p. 53);
2. Recalcitrance: if "the historical material resists being interpreted in terms of the initial judgments and preconceptions" (Ibid.), it forces us to revise our beliefs;
3. Confirmation: case studies can provide evidence for our philosophical beliefs in the sense that "the available evidence makes the belief in question more justified, better warranted, more plausible, more acceptable, or more likely to be true, than it would be if the corresponding evidence were not available" (Ibid.).
4. Adjudicating: "a case study may provide the philosophy of science with evidence that adjudicates between conflicting philosophical views" (Ibid).

Kinzel argues that historical case studies can provide evidence in the first three cases (though such evidence may still be of limited strength), but not in the fourth one. The reason why they cannot play a role of neutral adjudicators between rivaling philosophical views is primarily due to theory-ladenness of historical accounts (p. 55).

4.2 Introducing ABMs to HPS

We will now argue that ABMs can play all four of the above evidential functions, though just like in the case of historical case studies, their evidential import is limited.¹¹ This means not only that ABMs can play a legitimate epistemic role, but also that they can be used as a complementary method to historical case studies, mitigating some of the problems of the latter approach.

First, ABMs can provide *novel*, previously unknown and sometimes even surprising information. This includes providing novel explanatory hypotheses. For instance, ABMs may reveal social mechanisms underlying particular historical episodes, which we haven't been aware of. This is especially interesting in view of the fact that they are simulations: as such, they can provide insights into the evolution of historical episodes since we can observe socio-epistemic mechanisms "in motion". Zollman's ABMs are a case in point: they show how different social networks can lead to communities that are more or less efficient in acquiring knowledge. However, whether such explanatory hypotheses will in turn make the model explanatory of epistemically relevant phenomena (i.e. whether it will allow us to make relevant counterfactual inferences), depends on how successful we are in

¹¹ We are indebted to Erik Weber for the suggestion that ABMs could play the role of confirming historical hypotheses and explanations.

specifying the set of modeled explananda. As we have seen in the previous section, if a model offers an explanatory hypothesis for benefits of intensive alcohol consumption by ruling out numerous factors relevant for the assessment of the very same benefits, we could hardly call it explanatory in any epistemically relevant sense. Even though such a model would ‘explain’ a situation in which drinking alcohol all the time can be considered beneficial, it is hard to imagine a context in which such explanation would be of any epistemic interest. Hence, in order to convert an explanatory hypothesis of a model into its explanatory power, it is essential to determine its adequate target phenomenon.

Second, ABMs can play the role of *recalcitrance* by offering evidence that forces us to revise our beliefs concerning normative aspects of scientific inquiry. For instance, if Zollman effect were to hold under certain conditions that are relevant for actual scientific inquiry, we may have to revise our assumption that a high degree of connectedness among scientists is in most situations beneficial.

Third, ABMs can provide *confirmation* of our methodological and socio-epistemological assumptions in the same sense as explicated by Kinzel above: by raising the credibility and probability of our normative views.

Finally, ABMs can play a role in *adjudicating* between different normative philosophical conceptions about science by showing the superiority of one view over another (e.g. under specific conditions of scientific inquiry represented in the model). For instance, take the contemporary discussion on scientific pluralism as a normative methodological approach (see e.g. Kellert, Longino, and Waters 2006; Chang 2012) that is contrasted with a monist view on scientific methodology. In spite of attempts by pluralists to argue that the pluralist approach has numerous advantages and would lead to more efficient scientific inquiry,¹² concrete empirical arguments have been difficult to come by. Whether scientific inquiry proceeds more efficiently if the given scientific community divides the labor in such a way that multiple lines of inquiry are pursued at the same time, or whether they are more efficient if everyone focuses on only one of the rivaling accounts at the time – may be hard to answer from a philosophical armchair. Beside suffering from the problem of theory-ladenness, historical case studies may only provide single examples, rather than a general insight, sufficient to adjudicate this matter (as Chang (2012, 2015) himself remarks).

¹² Note however that beside arguing for a pluralist methodology in terms of efficiency of inquiry (or in terms of maximizing some other values at the level of the community), an argument for pluralism can also be given in view of ethical or political principles characterizing inquiry of individual scientists. In the latter case it suffices to show that a pluralist methodology fulfills such ethical or political criteria, irrespective of its efficiency (or maximization of values) at the level of the scientific community. Hence, our point that ABMs could be useful only concerns the former approach to pluralism, which argues for it in terms of its efficiency/value maximization at the level of the community.

This is where ABMs could perhaps be helpful: different simulations of scientific inquiry could provide evidence for the efficiency of different (monist and pluralist) methodologies under a variety of socio-epistemic circumstances. For instance, if we can establish that the superiority of a pluralist methodology is a robust property that holds across a wide range of parameters and changes in idealizing assumptions, this would be a strong argument in favor of pluralists. Alternatively, if it turns out that the matter is much more context dependent, we might be able to specify the contexts in which one methodological approach is superior to another. Of course, this should not be taken as suggesting that the construction of ABMs is not theory laden, nor that they can necessarily adjudicate between opposing methodological conceptions. While the robustness analysis discussed in Section 3 can help in disclosing assumptions underlying a given model, it could also happen that we end up without any robust property that holds across various models. In the latter case ABMs obviously wouldn't be of much help when it comes to adjudication. However, whether they can perform this function or not is an empirical question that would have to be examined for each philosophical dispute on its own.

Altogether, ABMs can perform each of the above four functions with respect to a certain target phenomenon *only* if the results obtained by them have been shown to be robust in two senses indicated in the previous section: under the changes in the relevant parameters and under the changes in idealizing assumptions. Since each of these robustness analyses may be a question without a positive test (but only a negative one), the evidence obtained by a given ABM will be as strong as detailed and exhaustive its robustness analysis is.

Even though Zollman initiated the research on ABMs that goes in this direction, his initial results, as we have seen, failed to satisfy robustness criteria, necessary for the above explicated evidential and explanatory roles. In the next section we will illustrate how we can build on his model, in order to bring it closer to an ABM that can actually provide evidence for the PUD case.

5 Towards an ABM of scientific interaction that is epistemically functional

As we have seen, the lack of robustness of Zollman's results under parameter changes has been shown by Rosenstock, O'Connor, and Bruner (2017). In view of this the authors suggest that the result may hold only in the context of difficult inquiry. In our (Frey and Šešelja 2018) we examine the robustness of Zollman's results under the changes in idealizing assumptions, relevant for the context of difficult inquiry. In this way we specify the context of (difficult) inquiry in which Zollman's findings are more likely to hold, and those in which they are less likely to hold. In this section we will show how our model can be used to examine the validity of Zollman's hypothesis (namely, the Zollman effect) for the concrete historical case study: the above mentioned research on PUD. We will first present the main features of our ABM and then turn to its application to this case study.

5.1 Our model

Our ABM is built by replacing a number of assumptions in Zollman's (2010) model. We will now present a brief version of these replacements:¹³

Dynamic epistemic success

First, instead of the assumption that the epistemic success of two rivaling theories is static (i.e. having fixed values for OPS, 0.5 and 0.499), we assume that scientists' success in finding the corroborating evidence for their theories changes over time. In particular, if they are pursuing the better theory, agents will be more and more successful, and if they are pursuing the worse theory they will be less and less successful. This feature is implemented by means of 'restless bandits', which are such that their OPS changes over time. To avoid ambiguities, we have replaced the notion of OPS with the notion of a scientist's *current probability of success* (CPS), assigned to each theory, representing the probability of gaining corroborating evidence for it, given the current state of one's inquiry. In addition, we call the value which an agent assigns to her current theory in view of Bayesian updating the *subjective probability of successes* (SPS) which she assigns to the given theory at the given time point. While at the beginning of a run the CPS of each theory is, respectively, 0.5 and 0.499, every X time steps the CPS of the better theory gradually increases towards 1, while the CPS of the worse theory gradually decreases towards 0.¹⁴ Our assumption is motivated by cases such as PUD, where successful applications of the bacterial hypothesis increased, while those of the acidity hypothesis decreased.¹⁵

An important corollary of our implementation of dynamic epistemic success is that due to the improvements in CPS of the theories, all our scientists will *eventually* discover which theory is objectively better. Hence, even if scientists prematurely abandon the better of the two hypotheses, they will eventually get back to it. This means that the question of the efficiency of inquiry shifts from 'How often are

¹³ The model is programmed in NetLogo (Wilensky 1999). For the technical details of the model see Frey and Šešelja 2018 as well as the open-source code, available on GitHub: <https://github.com/daimpi/SocNetABM/tree/RobIdeal>.

¹⁴ The CPS assigned by an agent to the given theory after an update, expressed in terms of the CPS of the same theory before the update will be:

$$CPS_{after}(T) = CPS_{before}(T) + f(d), \text{ where } f(d) = \frac{d}{1000} \text{ and } d = APS(T) - CPS_{before}(T), \text{ and}$$

where APS (actual probability of success) of the better theory is 1 and APS of the worse theory is 0.

¹⁵ Note that the CPS of each theory virtually never reach 1 and 0, respectively, since the runs end long before these values would be assigned. Moreover, the results presented in this section remain robust if we replace the values of 1 and 0 with 0.9 and 0.3, which correspond to the contemporary success rate in treating PUD with the antibiotic vs. antacid treatment (Hosking et al. 1994; Moayyedi et al. 2000).

scientists successful?’ to ‘How long do they take to get it right?’. Note that this doesn’t mean that consensus on the worse theory is not captured by the model. To the contrary, the scientific community may still abandon the better theory for a large time frame of the given run. As a result, such a community will be much slower in switching back and converging on the better theory. Hence, instead of measuring efficiency in terms of the percentage of successful runs (as Zollman does), we measure efficiency in terms of time that scientists need to converge on the better theory. This assumption nicely fits the PUD case, where scientists abandoned the bacterial hypothesis, only to get back to it 30 years later. Moreover, even though the two hypotheses appeared equally promising in the first half of the twentieth century, the contemporary antibiotic treatment (based on the bacterial hypothesis) exhibits the empirical success rates of around 90% vs. less than 30% for the antacid treatment (Hosking et al. 1994; Moayyedi et al. 2000), which are quite different values from the almost indistinguishable success rates for both theories employed in Zollman’s original model (namely, 0.5 and 0.499).

Hence, the notion of dynamic epistemic success is the basic assumption which we employ in all our simulations. The remainder of the assumptions concern factors that may be present or absent in our runs.

Critical interaction

Starting from the assumption of dynamic epistemic success of the pursued theories, we examine what happens if scientists interact critically. In order to test the impact of this feature of scientific communication, we replace Zollman’s assumption that all scientific interaction is epistemically equal with the assumption that scientists sometimes interact *critically*, where such interaction is epistemically beneficial. Looking at actual scientific inquiry, critical interaction plays an important role in disclosing errors that regularly appear in scientific research, namely false positives and false negatives. As such, critique tends to be truth conducive since it allows for the exposure of false beliefs. We implement this feature by assuming that criticism is triggered every time an agent pursuing theory T_x receives information from an agent pursuing T_y , such that T_y turns out to be better than she has expected.¹⁶

Similarly to the implementation of dynamic epistemic success, we implement critical interaction by allowing agents to slightly improve the CPS of their current theory towards 1 or 0 whenever they receive criticism, depending on whether they are pursuing the better or the worse theory.¹⁷

¹⁶ More precisely, scientist S_1 working on T_x is affected by criticism whenever the success rate of the rivaling theory (reported by scientist S_2 working on T_y) from the pulls in the most recent round is higher than the value S_1 has had in her memory, i.e. in case for S_1 : $SPS_{before}(T_y) < SPS_{after}(T_y)$.

¹⁷ Here we employ the same formula presented in Footnote [14].

Rational inertia

The next feature we examine is rational inertia that scientists may have towards their current theories. While Zollman's agents are easily swayed by new evidence, in the context of actual scientific inquiry scientists tend to retain their current theory at least for a certain time period. More precisely, they tend to stick with it unless they are convinced that it can no longer be saved from the defeating evidence. This is not irrational behavior: if a scientist knows that her current evidence is insufficient to determine whether the theory could eventually be accepted—as it may easily occur in times of difficult inquiry—it would be irrational to abandon it before attempting its further development, and rational to stick to it for a while longer (see (Kelp and Douven 2012)). In addition, changing one's inquiry usually includes a number of costs (e.g. acquiring additional knowledge, new equipment, etc.), which is another motivation for such inertia.

We implement this feature in terms of a *jump threshold*: agents 'jump' to the rivaling theory only after the latter has according to one's beliefs (i.e. in view of one's SPS) turned out to be better than their current theory for Y number of rounds. For the sake of simplicity, we will denote 'jump threshold with the value of Y ' by 'inertia of Y '.

Equally promising rivals

Finally, we examine what happens if we assume that there is an interval within which rivaling theories count as equally good, which replaces Zollman's assumption that one always has a linear preference order over the rivaling theories. We implement this feature in terms of the assumption that the rival theory counts as better only if it surpasses one's own theory by the margin of 0.1 (in terms of SPS that an agent assigns to the theories).

The results

The results of (Frey and Šešelja 2018)¹⁸ suggest that Zollman's finding concerning the superiority of the cycle towards the complete graph holds only if it is neither the case that scientists have rational inertia, nor is it the case that they treat rivaling theories as equally good within a certain interval. As soon as one of these two conditions is satisfied, our ABM indicates that the cycle ceases to be superior to the complete graph, or it even becomes inferior to it.

¹⁸ The simulations were run for up to 100,000 rounds, for populations consisting of 4 to 11 scientists, and we recorded the point when agents converged on the better theory (by convergence on a theory we mean that all agents end up on that theory, without switching back to the rival). Each scenario is based on 10,000 runs. For the changes in CPS we have opted for (minor) changes every 100 rounds. Our results remain stable if the rate at which CPS changes is every 10 rounds, or every 500 rounds. For rational inertia we have chosen the value of 10.

5.2 Examining the case of PUD

But what about the case of PUD – where does it fall within the above spectrum of assumptions? In order to answer this question, we need to do some empirical calibration of our model.

First of all, we will extract results of those runs (for different combinations of the above mentioned factors) in which both theories are populated at the beginning of the run. Since the distribution of agents at the beginning of the run is random, it could happen that there are no agents on one of the theories. In the case of PUD research, however, we know that there were scientists pursuing both theories from the second half of the nineteenth century onward (Šešelja and Straßer 2014).

The analysis of our results shows a general correlation between the time scientists spend exclusively on the wrong theory and the time they need to reach a convergence on the better theory. This is not surprising: the longer a community keeps a consensus on the worse theory, the longer it will need to switch back to the better one. In other words, the longer a community stays on the wrong theory, the less efficient it is.¹⁹ Figure [2] shows the average time scientists connected in the cycle vs. the complete graph spend on the wrong theory for different scenarios. These results suggest that the cycle performs better than the complete graph in each of the given scenarios. Note that we have omitted the scenario in which scientists treat rivaling theories as equally good within the interval of 0.1 due to the fact that in these runs scientists don't exhibit consensus on the wrong theory, and hence those runs are not representative of the PUD case (we will come back to this point below, concerning the normative conclusions that can be drawn from our model). It is also interesting to notice that changing the network type (from the cycle to the complete graph) may not increase the efficiency of the community as much as, for example, introducing caution to theory assessment (represented in terms of inertia).

¹⁹ Note though that the duration of a consensus on the wrong theory is not the only factor determining how long the scientists need to converge on the right theory. Another relevant factor is, e.g. the time needed for a group to converge on the right theory, after the consensus on the wrong one has been abandoned.

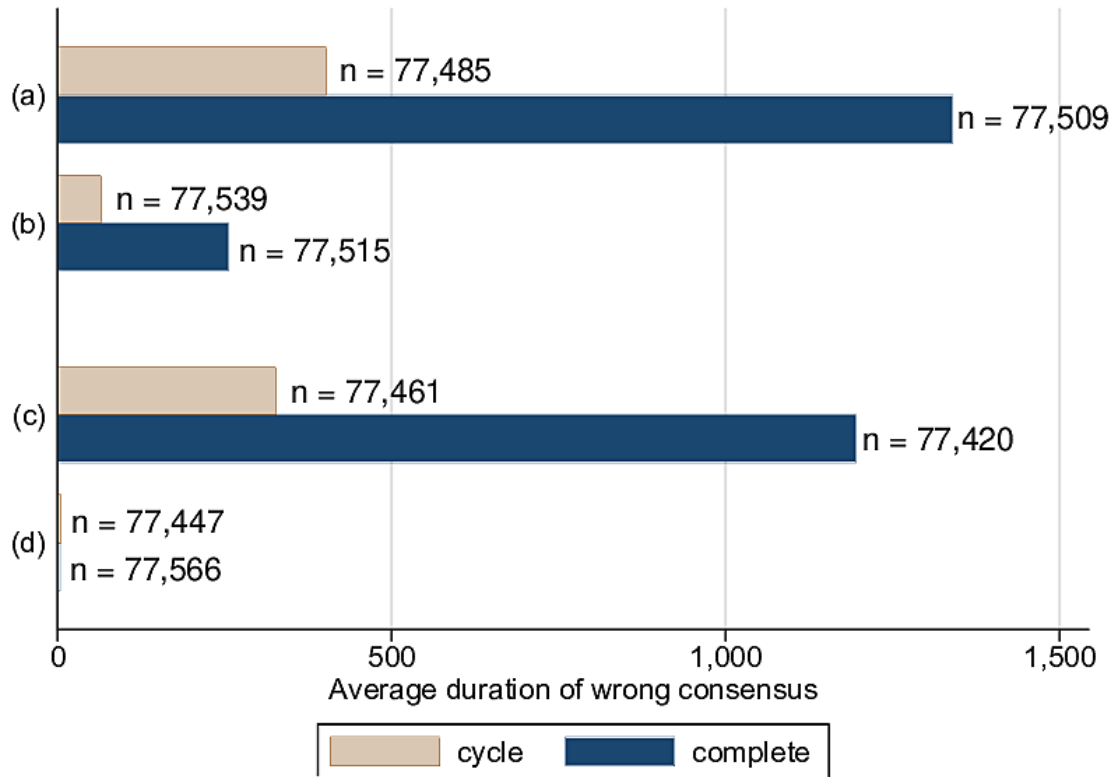


Figure 2: Average number of rounds a consensus on the worse theory lasted aggregated over all population sizes (4 - 11) for different treatments: (a): no critical interaction, no inertia; (b): no critical interaction, inertia of 10; (c): critical interaction, no inertia; (d): critical interaction, inertia of 10 (though the time for this treatment is so short that the bars are hardly visible, the complete graph here needs less time than the cycle); n stands for the number of runs falling into the given category.

Nevertheless, these results don't tell us yet much about the PUD case itself. For this question, we need to make some further calibrations. One source of information that is uncontroversial is the timeline of this inquiry. We know that the research on PUD started in the late nineteenth century (or earlier), that the two theories were pursued until the 1950s, and that there was a wide-spread consensus on the acidity hypothesis until the 1980s. Finally, we have a successful consensus on the bacterial theory in the mid-1990s.²⁰ Altogether, we can specify the portion of the overall time

²⁰ Actual historical timeline was a bit more complex: for example, the acidity hypothesis wasn't merely dismissed. Rather, the antacid treatment was retained in those cases in which the PUD is not caused by *Helicobacter pylori*. Given the idealized setup of the model—namely, the idea that converging on only one of the rivals counts towards success of the scientists in finding the right theory—we will stick with the above idealized timeline in order to be able to represent the return to the bacterial hypothesis as a successful outcome.

of this inquiry during which the scientific community pursued only the acidity hypothesis.

If T stands for the overall time of the run and C stands for the time during which the scientific community had a consensus on one of the theories, then assuming that $T > 90$ (i.e. assuming the inquiry started in 1900 or earlier and ended in the 1990s) and $C = 30$ (since the scientific community pursued only the acidity hypothesis from the 1950s to the 1980s), we are interested only in those runs for which $T/C > 90/30$ i.e. $T/C > 3$.²¹

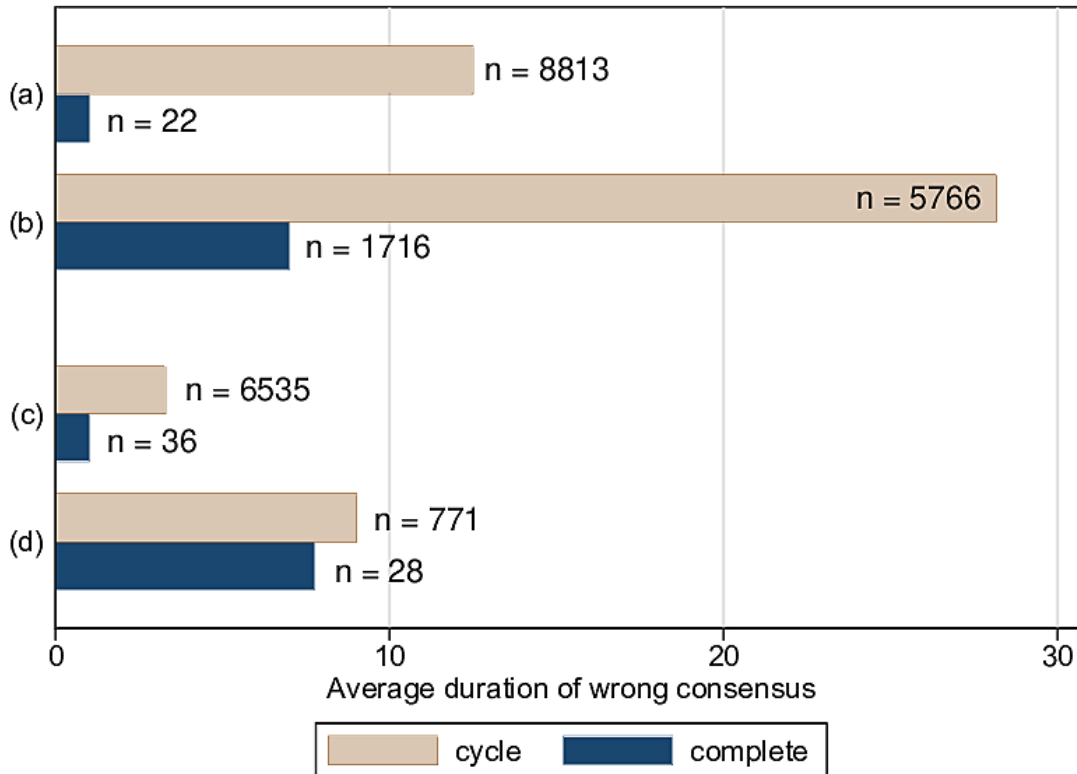


Figure 3: Average number of rounds a consensus on the worse theory lasted aggregated over all population sizes but limited to runs which exhibited consensus for less than 1/3 of the run. For the meaning of shortcuts used for the treatments see the caption above in Figure 2; n stands for the number of runs falling into the given category.

Filtering runs that satisfy this criterion produces results shown in Figure [3]. Here the cycle seems more likely to produce runs in which scientists are stuck on the wrong path for a longer time period than the complete graph. It is important to note

²¹ Note that the absolute values for T and C are not to be taken as representative of actual time periods since we haven't established the mapping between the time steps in the model and the real time. This is why we only use the relative value (T/C), indicating that for two thirds of the overall time of inquiry diversity is preserved.

though that the process of filtering drastically reduces the number of runs that fall into this category (see Figure [4]): out of the 309,932 cycle runs there are only 21,705 left, while out of the 310,010 complete graph runs even fewer meet this criterion: only 1,802 runs. The reason why the vast majority of runs with the complete graph are filtered out is that the ratio T/C tends to be much smaller in this case, since scientists tend to reach a consensual pursuit of only one of the theories much more often. As a result, cases like PUD, where for the large portion of inquiry scientists pursued both theories, is hard to find among the runs with fully connected communities.

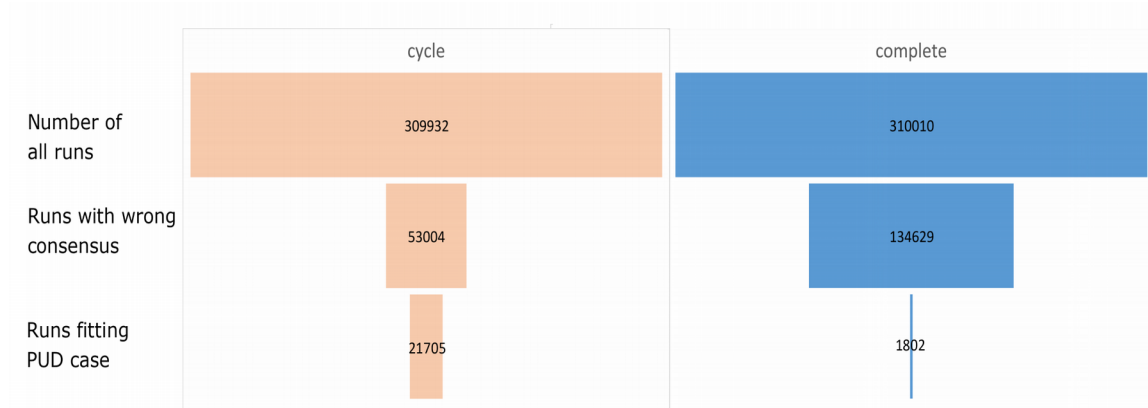


Figure 4: Number of runs aggregated over all simulations in which both theories are populated at the beginning of the run. ‘Runs with wrong consensus’ are those in which scientists at some point during the run all work on the worse theory. ‘Runs fitting PUD case’ are those for which $T/C > 3$.

This is further evidence that, even though the wide information sharing may seem to be the cause of problems in some contexts of scientific inquiry, it doesn’t seem to be the likely cause of the abandonment of the bacterial hypothesis in the PUD case. Rather, the culprit seems more likely to lie either in the lack of critical interaction, or in the lack of caution expressed in terms of an interval within which rivaling theories count as equally good. The latter point concerns the runs in which we employ an interval of 0.1 within which theories count as equally promising and where scientists never exhibit consensus on the wrong theory (and which are thus absent from Figure [3]).

Nevertheless, our analysis is just the first step in the calibration of this model towards the PUD case, and hence towards the robustness analysis of our model aimed at providing evidence for this historical episode. In order to obtain reliable evidence, we need to map further parameters in the model into the space of the real-world phenomena. This way we can identify the portion of the parameter space that is relevant for the PUD case. Such a mapping would include, for instance, the size of the scientific community investigating PUD at the time²², time steps in the model in

²² Even though the population size in our model goes up to 11 agents, this could also be understood as representing 11 different research groups, working on PUD.

terms of time of actual scientific inquiry, the impact of extreme prior beliefs (as discussed originally by Zollman (2010)), etc.

Finally, it remains a task for future research to examine the robustness of this result in view of alternative models of scientific interaction. While an advantage of our current ABM lies in its specific representation of a dynamic epistemic success, which allows us to mirror the development of the PUD case, other ABMs could easily be adjusted to this end as well. A suitable candidate is our argumentative ABM presented in (Borg et al. 2018), which shares a number of features with the current model (e.g. the presence of critical interaction, rational inertia, the interval within which theories count as equally good, etc.), but which is nevertheless based on different assumptions concerning the representation of scientific inquiry (e.g. knowledge acquisition and scientific interaction are represented in terms of arguments and argumentative exchange). It will be interesting to see whether incautious theory assessment, which according to our current results appears to be the main culprit in the PUD case, will turn out to be a robust property, resulting from the above mentioned argumentative ABM as well. If the opposite turns out to be the case (i.e. if some other factors turn out to be a more likely culprit) this will indicate that the two models represent different target phenomena, the specification of which would be yet another question worthy of pursuit.

6 Conclusion

In this paper we have examined the epistemic function of highly idealized ABMs of scientific inquiry. Focusing on ABMs of scientific interaction we have suggested that they are currently only of exploratory, rather than explanatory use. Next, following Kinzel's (2015) discussion on the epistemic function of historical case studies, we have argued that ABMs can play both evidential and explanatory roles, under the condition that they have passed adequate robustness analyses. We have illustrated our point with an ABM of scientific interaction, built on the basis of Zollman's (2010) model and aimed at explaining the dynamics underlying the research on PUD.

Our results can be understood as playing the roles of novelty and recalcitrance: on the one hand, by showing that (and explaining why) the complete graph doesn't seem likely to represent the dynamics of the PUD episode, and on the other hand, by triggering a revision of the belief obtained by means of Zollman's model, that too wide information sharing was a likely causal mechanism leading scientists to abandon the bacterial hypothesis. Whether our current conclusion will have to be revised yet again, in view of further robustness analyses, remains a question for future research.

A more general take-home message of our paper was showing the significance of empirical calibration of ABMs in case we should consider them explanatory of their

intended target phenomena.²³ Given a recent trend of simple ABMs of science proposed in the literature, this task is all the more important. For each of these models, whether they concern network effects or, for instance, the division of cognitive labor (see Section 1), it is essential to examine whether their results are robust once a given model is adjusted to represent a concrete case – its purported target. To this end, it remains a task for future research to specify benchmark case studies for each individual model, which can serve as the basis for their respective robustness checks.²⁴

Authors

Daniel Frey (freyda87@gmail.com) is a master student of economics at the Faculty of Economics and Social Sciences, Heidelberg University, with interests in social epistemology and philosophy of science. He is an associate member of the Research Group for Non-Monotonic Logics and Formal Argumentation at Ruhr-University Bochum, with whose members he has been collaborating in the domain of agent-based models of scientific inquiry.

Dunja Šešelja (dunja.seselja@rub.de) is a postdoctoral researcher at the Institute for Philosophy II, Ruhr-University Bochum, at the Chair for Logic and Epistemology, and she is an associate member of the Research Group for Non-Monotonic Logics and Formal Argumentation. Her research lies at the intersection of integrated history and philosophy of science, methodology of science, social epistemology, and formal approaches to scientific inquiry. Her recent research focuses on agent-based models of scientific inquiry and their epistemic function.

REFERENCES

Alexander, Jason McKenzie, Johannes Himmelreich, and Christopher Thompson. 2015. "Epistemic Landscapes, Optimal Search, and the Division of Cognitive Labor." *Philosophy of Science* 82 (3). JSTOR:424–53.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbours." *The Review of Economic Studies* 65 (3). Wiley-Blackwell:595–621.

Bokulich, Alisa. 2014. "How the Tiger Bush Got Its Stripes: 'How Possibly' vs. 'how Actually' model Explanations." *The Monist* 97 (3). Oxford University Press Oxford, UK:321–38.

Borg, AnneMarie, Daniel Frey, Dunja Šešelja, and Christian Straßer. 2017a. "An Argumentative Agent-Based Model of Scientific Inquiry." In *Advances in Artificial Intelligence: From Theory to Practice: 30th International Conference on Industrial*

²³ This is in line with the emphasis on empirical calibration by Martini and Pinto (2016).

²⁴ We are grateful to two anonymous reviewers for their valuable comments and suggestions on the earlier draft of this paper.

Engineering and Other Applications of Applied Intelligent Systems, Iea/Aie 2017, Arras, France, June 27-30, 2017, Proceedings, Part I, edited by Salem Benferhat, Karim Tabia, and Moonis Ali, 507–10. Cham: Springer International Publishing.

———. 2017b. “Examining Network Effects in an Argumentative Agent-Based Model of Scientific Inquiry.” In *Logic, Rationality, and Interaction: 6th International Workshop, Lori 2017, Sapporo, Japan, September 11-14, 2017, Proceedings*, edited by Alexandru Baltag, Jeremy Seligman, and Tomoyuki Yamada, 391–406. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-55665-8_27.

———. 2018. “Epistemic Effects of Scientific Interaction: Approaching the Question with an Argumentative Agent-Based Model.” *Historical Social Research*.

Chang, Hasok. 2012. *Is Water H₂O? Evidence, Pluralism and Realism*. Springer.

———. 2015. “The Chemical Revolution Revisited.” *Studies in History and Philosophy of Science Part A* 49. Elsevier:91–98.

De Langhe, Rogier. 2013. “Peer Disagreement Under Multiple Epistemic Systems.” *Synthese* 190. Springer:2547–56.

———. 2014. “A Unified Model of the Division of Cognitive Labor.” *Philosophy of Science* 81 (3). JSTOR:444–59.

Douven, Igor. 2009. “Uniqueness Revisited.” *American Philosophical Quarterly* 46 (4). JSTOR:347–61.

———. 2010. “Simulating Peer Disagreements.” *Studies in History and Philosophy of Science Part A* 41 (2). Elsevier:148–57.

Dray, William H. 1957. “Laws and Explanation in History.” Oxford: Oxford University Press.

Forber, Patrick. 2010. “Confirmation and Explaining How Possible.” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 41 (1). Elsevier:32–40.

———. 2012. “Conjecture and Explanation: A Reply to Reydon.” *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 43 (1). Elsevier:298–301.

Frey, Daniel, and Dunja Šešelja. 2018. “Robustness and Idealization in Agent-Based Models of Scientific Interaction.” *British Journal for the Philosophy of Science*, *Conditionally Accepted*.

Gelfert, Axel. 2016. *How to Do Science with Models: A Philosophical Primer*. Springer.

Goldman, Alvin, and Thomas Blanchard. 2016. “Social Epistemology.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2016.

<http://plato.stanford.edu/archives/sum2016/entries/epistemology-social/>;
Stanford University.

Grim, Patrick. 2009. "Threshold Phenomena in Epistemic Networks." In *AAAI Fall Symposium: Complex Adaptive Systems and the Threshold Effect*, 53–60.

Grim, Patrick, Daniel J Singer, Steven Fisher, Aaron Bramson, William J Berger, Christopher Reade, Carissa Flocken, and Adam Sales. 2013. "Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success, and Convergence." *Episteme* 10 (04). Cambridge Univ Press:441–64.

Grüne-Yanoff, Till. 2009. "Learning from Minimal Economic Models." *Erkenntnis* 70 (1). Springer:81–99.

Hosking, S.W., S.C.S. Chung, M.Y. Yung, A.K.C. Li, J.J.Y. Sung, T.K.W. Ling, and A.F.B. Cheng. 1994. "Duodenal Ulcer Healing by Eradication of Helicobacter Pylori Without Anti-Acid Treatment: Randomised Controlled Trial." *The Lancet* 343 (8896):508–10. [https://doi.org/http://dx.doi.org/10.1016/S0140-6736\(94\)91460-5](https://doi.org/http://dx.doi.org/10.1016/S0140-6736(94)91460-5).

Kellert, Stephen H., Helen E. Longino, and C. Kenneth Waters, eds. 2006. *Scientific Pluralism*. Minneapolis: University of Minnesota Press.

Kelp, C., and I. Douven. 2012. "Sustaining a Rational Disagreement." *EPSA Philosophy of Science: Amsterdam 2009*. Springer, 101–10.

Kinzel, Katherina. 2015. "Narrative and Evidence. How Can Case Studies from the History of Science Support Claims in the Philosophy of Science?" *Studies in History and Philosophy of Science Part A* 49. Elsevier:48–57.

Kummerfeld, Erich, and Kevin JS Zollman. 2016. "Conservatism and the Scientific State of Nature." *The British Journal for the Philosophy of Science* 67 (4). Br Soc Philosophy Sci:1057–76.

Kuorikoski, Jaakko, and Aki Lehtinen. 2009. "Incredible Worlds, Credible Results." *Erkenntnis* 70 (1). Springer:119–31.

Lehtinen, Aki. 2017. "Derivational Robustness and Indirect Confirmation." *Erkenntnis*. Springer, 1–38.

Lehtinen, Aki, and Jaakko Kuorikoski. 2007. "Unrealistic Assumptions in Rational Choice Theory." *Philosophy of the Social Sciences* 37 (2). Sage Publications Sage CA: Los Angeles, CA:115–38.

Martini, Carlo, and Manuela Fernández Pinto. 2016. "Modeling the Social Organization of Science." *European Journal for Philosophy of Science*. Springer, 1–18.

Mäki, Uskali. 2009. "Realistic Realism About Unrealistic Models." In *The Oxford Handbook of Philosophy of Economics*, edited by Harold Kincaid and Don Ross, 68–98. Oxford University Press.

- Moayyedi, Paul, Shelly Soo, Jonathan Deeks, David Forman, James Mason, Michael Innes, and Brendan Delaney. 2000. "Systematic Review and Economic Evaluation of Helicobacter Pylori Eradication Treatment for Non-Ulcer Dyspepsia." *BMJ* 321 (7262). BMJ Publishing Group Ltd:659–64.
<https://doi.org/10.1136/bmj.321.7262.659>.
- Pfleiderer, Paul. 2014. "Chameleons: The Misuse of Theoretical Models in Finance and Economics." *Revista de Economía Institucional* 16 (31). Revista de Economía Institucional:23–60.
- Pitt, Joseph C. 2001. "The Dilemma of Case Studies: Toward a Heraclitian Philosophy of Science." *Perspectives on Science* 9 (4). MIT Press:373–82.
- Pöyhönen, Samuli. 2017. "Value of Cognitive Diversity in Science." *Synthese* 194 (11). Springer:4519–40.
- Reutlinger, Alexander, Dominik Hangleiter, and Stephan Hartmann. 2016. "Understanding (with) Toy Models." *The British Journal for the Philosophy of Science*. Oxford University Press.
- Reydon, Thomas AC. 2012. "How-Possibly Explanations as Genuine Explanations and Helpful Heuristics: A Comment on Forber." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 43 (1). Elsevier:302–10.
- Riel, Raphael van. 2015. "The Content of Model-Based Information." *Synthese* 192 (12). Springer:3839–58.
- Rosenstock, Sarita, Cailin O'Connor, and Justin Bruner. 2017. "In Epistemic Networks, Is Less Really More?" *Philosophy of Science* 84 (2):234–52.
- Soler, Léna, Emiliano Trizio, Thomas Nickles, and William Wimsatt. 2012. *Characterizing the Robustness of Science: After the Practice Turn in Philosophy of Science*. Vol. 292. Springer.
- Strevens, Michael. 2010. "Reconsidering Authority: Scientific Expertise, Bounded Rationality, and Epistemic Backtracking." *Oxford Studies in Epistemology* 3. OUP Oxford:294.
- Šešelja, Dunja, and Christian Straßer. 2014. "Heuristic Reevaluation of the Bacterial Hypothesis of Peptic Ulcer Disease in the 1950s." *Acta Biotheoretica* 62:429–54.
- Thiele, Jan C, Winfried Kurth, and Volker Grimm. 2014. "Facilitating Parameter Estimation and Sensitivity Analysis of Agent-Based Models: A Cookbook Using Netlogo and R." *Journal of Artificial Societies and Social Simulation* 17 (3):11.
- Thoma, Johanna. 2015. "The Epistemic Division of Labor Revisited." *Philosophy of Science* 82 (3). JSTOR:454–72.

- Weisberg, Michael, and Ryan Muldoon. 2009. "Epistemic Landscapes and the Division of Cognitive Labor." *Philosophy of Science* 76 (2). JSTOR:225–52.
- Wilensky, Uri. 1999. "NetLogo.(<http://Ccl.northwestern.edu/Netlogo/>)." *Center for Connected Learning and Computer Based Modeling, Northwestern University*.
- Wray, K Brad. 2011. *Kuhn's Evolutionary Social Epistemology*. Cambridge University Press.
- Ylikoski, Petri, and N Emrah Aydinonat. 2014. "Understanding with Theoretical Models." *Journal of Economic Methodology* 21 (1). Taylor & Francis:19–36.
- Zollman, Kevin J. S. 2007. "The Communication Structure of Epistemic Communities." *Philosophy of Science* 74 (5). JSTOR:574–87.
- . 2010. "The Epistemic Benefit of Transient Diversity." *Erkenntnis* 72 (1). Springer:17–35.
- . 2013. "Network Epistemology: Communication in Epistemic Communities." *Philosophy Compass* 8 (1). Wiley Online Library:15–27.