The Cultural Evolution of Methods in Philosophy of Science: Model & Data

Rafael Ventura

Abstract

What is the relation between philosophy of science and the sciences? As Pradeu et al. (2021) and Khelfaoui et al. (2021) recently show, part of this relation is constituted by "philosophy in science": the use of philosophical methods to address questions in the sciences. But another part is what one might call "science in philosophy": the use of methods drawn from the sciences to tackle philosophical questions. In this paper, we focus on one class of such methods and examine the role that model-based methods play within "science in philosophy". To do this, we first build a bibliographic coupling network with Web of Science records of all papers published in philosophy of science journals from 2000 to 2020 ($N = 9,217$). After detecting the most prominent communities of papers in the network, we use a supervised classifier to identify all papers that use model-based methods. Drawing on work in cultural evolution, we also propose a model to represent the evolution of methods in each one of these communities. Finally, we measure the strength of cultural selection for model-based methods during the given time period by integrating model and data. Results indicate not only that model-based methods have had a significant presence in philosophy of science over the last two decades, but also that there is considerable variation in their use across communities. Results further indicate that some communities have experienced strong selection for the use of model-based methods but that other have not; we validate this finding with a logistic regression of paper methodology on publication year. We conclude by discussing some implications of our findings and suggest that model-based methods play an increasingly important role within "science in philosophy" in some but not all areas of philosophy of science.

1 Introduction

What is the relation between philosophy of science and the sciences? To answer this question, philosophers of science have recently turned to digital techniques and bibliometric data (Malaterre et al., 2019, 2020; Khelfaoui et al., 2021; Pradeu et al., 2021). This approach has made it possible to identify philosophers who regularly publish in science journals, track how often philosophy papers are cited by scientists, and measure the impact that philosophers have within scientific disciplines. A main finding from this emerging body of work is that philosophers often make genuine contributions to scientific debates by relying on methods that are typically regarded as philosophical, such as conceptual analysis and metaphysical theorizing. This is what some now call "philosophy in science" (Khelfaoui et al., 2021; Pradeu et al., 2021).

However, the relation between philosophy of science and the sciences is not unidirectional. Although "philosophy in science" is certainly part of the picture, philosophers of science also engage with the sciences by drawing on methods from scientific disciplines to address philosophical questions. There are of course many different ways in which this can occur. For instance, philosophers of science sometimes borrow survey-based and experimental methods from the cognitive and behavioral sciences in what is now known as "experimental philosophy of science" (Knobe, 2007; Griffiths and Stotz, 2008; Machery, 2016). Philosophers of science can also address philosophical questions by relying on digital tools to analyze bibliometric data (Pence and Ramsey, 2018; Ramsey and De Block, 2021), as in the studies described above. Another class of methods that philosophers of science can and often do borrow from the sciences are model-based methods, such as mathematical and computational models (Wheeler, 2013; Leitgeb, 2013; Mayo-Wilson and Zollman, 2021). Thus, another side of the relation between philosophy of science and the sciences is what one might call "science in philosophy": the use of methods drawn from the sciences to tackle philosophical questions.

Although surveys, experiments, tools for bibliometric data analysis, and models are widely used in the sciences, they make up a very heterogeneous collection of methods. It is therefore challenging to study their use in philosophy of science at once, especially when relying on the automated tools we describe below. For these reasons, we focus here on the use of a single type of method: model-based methods. Model-based methods make up a complex class of methods that has sparked a large and growing philosophical literature (Suárez, 2008; Weisberg, 2012; Frigg et al., 2020). Our goal here is not to contribute to our understanding of how models are used in science. Rather, it is to understand how philosophers borrow model-based methods from the sciences to address question in philosophy. The use of such methods is especially common in philosophy of science. A recent example is Sprenger and Hartmann (2019), who make extensive use of probability theory to model scientific reasoning and address long-standing issues in general philosophy of science. Or take subfields of philosophy of science, such as philosophy of physics and philosophy of biology. In these subdisciplines, differential geometry and dynamical systems theory are important tools for building models, as recent work by Huggett and Wüthrich (2018) and Tanaka et al. (2020) illustrate. In work on the social dimension of science, numerical techniques such as computer simulations and agent-based models are quite widespread as well—for an early and a recent example, see Zollman (2007) and Weatherall et al. (2020).

In this paper, we examine the role that model-based methods play within "science in philosophy". To do so, we analyze a large bibliometric dataset. Using publicly available data from the Web of Science, we build a bibliographic coupling network with all research articles published in the main philosophy of science journals from 2000 to 2020. After detecting the most prominent communities of papers in the network, we use a supervised classifier to identify the papers that use model-based methods. Drawing on work in cultural evolutionary theory (Zollman, 2018; O'Connor, 2019; Heesen, 2019), we also propose a model to represent the evolution of methods in philosophy of science during the time period. By integrating this model with bibliometric data, we measure the strength of cultural selection for the use of model-based methods in philosophy of science. This allows us to not only determine the prevalence of model-based methods in philosophy of science, but also to test the hypothesis that there has been selection for the use of such methods.

Our results indicate that model-based models have had a significant presence in philosophy of science over the last two decades. We also find that there is considerable variation in the use of model-based methods in philosophy of science across different communities: while model-driven techniques are widespread in some, models are almost entirely absent from others. Moreover, we find that some communities have experienced strong selection for the use of model-based methods but that others have not. Our results therefore suggest that model-based methods play an increasingly important role in some but not all areas of philosophy of science.

The paper proceeds as follows. In Section 2, we present our data, describe the methods we use to analyze it, and introduce a model to represent the cultural evolution of methods in the philosophy of science; technical details of these methods are described in the corresponding Appendices. In Section 3, we report our findings on the prevalence of model-based methods in philosophy of science. We show that some areas of philosophy of science have experienced strong selection for the use of such methods and validate these results with a logistic regression of paper methodology on publication year. In Section 4, we discuss some implications of our results for recent work by Fletcher et al. (2021), Khelfaoui et al. (2021), and Pradeu et al. (2021). In Section 5, we conclude by noting some limitations of our approach and suggesting a few directions for future studies.

2 Data & Model

To study the use of model-based methods in philosophy of science, we first collected data from the Web of Science (`www.webofscience.com`). Among other services, the Web of Science website provides an online database with detailed information on papers published in academic journals. Records generally contain information on paper title, abstract, authors, and cited references. For this study, we used the advanced search tool to extract full records for all papers published in the main philosophy of science journals. Included in this study were the nine journals in general philosophy of science already studied by Pradeu et al. (2021)—for a complete list of journal titles, see Table 1. We then manually downloaded and saved the $11,030$ full records matching our search criteria for the time period between 2000 and 2020. The search was restricted to this time period because older records often lack data such as abstract or cited references.

Table 1: List of journals, together with number of papers published in each journal (N). Considered were all journals in general philosophy of science studied by Pradeu et al. (2021).

	Journal Title	N
1	*British Journal for the Philosophy of Science*	965
2	*Erkenntnis*	1,535
3	*European Journal for the Philosophy of Science*	300
4	*Foundations of Science*	552
5	*International Studies in Philosophy of Science*	325
6	*Journal for General Philosophy of Science*	416
7	*Philosophy of Science*	1,824
8	*Studies in History & Philosophy of Science*	1,196
9	*Synthese*	3,917

Contained in this initial sample were not only research papers, but also reviews, obituaries, and other editorial materials. To limit our study to research articles in philosophy and facilitate analysis, records not tagged as research articles as well as record written in languages other than English were removed; records with a missing abstract or with missing references were also excluded.

With the $N = 9,217$ remaining papers, we built a bibliographic coupling network (Kessler, 1963). Bibliographic coupling networks take the similarity between two papers to be a function of how often they cite the same papers. In a bibliographic coupling network, a node therefore represents a paper and a link between two nodes represents the extent to which two papers cite the same references. In other words, a link represents the similarity between two papers with respect to the references that they cite. Bibliographic networks are therefore built on the assumption that papers sharing many unique references are likely to address similar questions, while papers that do not share many unique references are likely to engage with different topics—for a recent use of a bibliographic coupling network in philosophy, see Noichl (2021).

To build a bibliographic coupling network, we calculated the term frequency and the inverse-document frequency of references for each paper—for technical details on how to build a bibliographic coupling network, see Appendix 1: Bibliographic Coupling Network. The term frequency measures the importance that a particular reference has to a paper; the inverse-document frequency measures the importance of a particular reference to the entire corpus. We then combined the term frequency and the inverse-document frequency to obtain the $tfidf(p_i)$ score for each paper. The $tfidf$ score measures not only how important a particular reference is to a paper, but also how important the reference is to the entire corpus: it characterizes each paper in terms of the importance that each reference in the entire corpus has to the paper.

As already noted, a link between two papers in a bibliographic coupling network represents how similar they are with respect to the references that they cite. To build such a network, we therefore need to measure the similarity between every pair of papers. To do so, we used the cosine similarity between the $tfidf$ scores of each pair of papers. Although other measures of similarity between pairs of papers are in principle possible, the cosine similarity is a common measure of similarity between $tfidf$ scores. The cosine similarity thus serves as a proxy for how much each pair of papers engage the same research questions, ranging in the unit interval and with 0 denoting complete dissimilarity and 1 denoting complete similarity.

Upon building the bibliographic coupling network, we proceeded to detect communities of papers that engage similar research questions. There are of course many different methods to detect communities in a network. A simple, computationally efficient, and widely used one is the algorithm for community detection due to Blondel et al. (2008). This method finds discrete communities in a network by maximizing network modularity. Modularity is a measure of how well-connected nodes are to other nodes within the same community and how poorly connected nodes are to other nodes outside the same community. As links between nodes in a bibliographic coupling network represent similarity between papers, this algorithm detects communities by finding a partition of the network that maximizes how similar papers are to other papers within the same community but dissimilar to papers in other communities—for technical details on how to detect communities, see Appendix 2: Community Detection.

Having detected communities of papers in the network, we then used a naive Bayes classifier to label papers with respect to their methodology. Naive Bayes classifiers are a family of simple and computationally efficient classification algorithms that generally perform well in text classification (McCallum et al., 1998; Chandrasekar and Qian, 2016); as we report below, the naive Bayes classifier we used also performed quite well. Naive Bayes classifiers assign items to classes on the basis of features that items have. In particular, naive Bayes classifiers assign items to classes by assuming that the occurrence of a given feature in the set of all items is probabilistically independent from the occurrence of one another feature (hence the epithet "naive"). To assign an item to a particular class, naive Bayes classifiers first calculate the probability that the item belongs to different classes given the features that the item has and then assign the item to the class with the highest probability conditional on the features of the item.

In our case, we used a multinomial naive Bayes classifier to classify papers with respect to their methodology given the words occurring in their abstracts and the last name of the authors in their cited references. This means that items correspond to papers, classes correspond to the two types of methods that a paper might use (model-based method vs. no model-based method), and features correspond to words contained in a paper's abstract as well as the last name of the authors cited in the paper's reference section. In a multinomial naive Bayes classifier, features correspond to the number of times that a word appears in a paper's abstract and the number of times that a last name appears in a paper's reference section. Our naive Bayes classifier therefore assigns the label "uses a model-based method" or "does not use a model-based method" to a paper depending on the words that appear in the paper's abstract and the last name of the authors that the paper cites—for technical details on the naive Bayes classifier we used, see Appendix 3: Naive Bayes Classifier.

But to assign an item to a class, a naive Bayes classifier must first estimate the parameters that allow it to calculate the conditional probability that an item belongs to different classes, given its features. This means that a naive Bayes classifier must first be fed the conditional probability of features given different classes, the unconditional probability of features, and the unconditional probability of classes. As this is a supervised algorithm, a naive Bayes classifier must therefore rely on humans to provide it with a dataset of items, their features, and the classes that these items belong to in order to estimate parameters and assign new items to the classes of interest.

To estimate parameters, we randomly selected 500 papers from the set of $N = 9,217$ research papers written in English for manual labelling. Papers were labeled as using model-based methods or not using such methods. Out of 500 papers, 62 were found to use model-based methods; the full list of manually labelled papers is available in the repository provided below. Labeling was done according to the following rubric. First, we checked for the occurrence of any mathematical expressions or figures that might indicate the use of model-based methods. Second, we read the paper abstract to determine whether the paper used mathematical expressions or figures as an example, to provide a philosophical interpretation of models built by others, to extend or adapt previous models, or to built its own model. Papers were labeled as using model-based methods if they used probability theory, dynamical systems theory, differential geometry, or numerical and computational techniques to extend, adapt, or build a model. The choice to focus on these mathematical tools and techniques in particular was made on the basis of expert interviews with practicing philosophers of science working in a wide range of subdisciplines, including philosophy of biology, cognitive science, computer science, decision and game theory, physics, and social science. Papers were labeled as not using model-based methods if they did not use any of these methods, or if they used any of these methods as an example or to provide a philosophical interpretation of models built by others. When we could not determine this on the basis of the abstract alone, we read the full paper. Although one might conjecture that not all papers included here address philosophical questions, we take the fact that a paper was published in a philosophy journal as a proxy for it addressing philosophical questions.

In addition to facilitating replication, this rubric serves an important function: it allows us to distinguish papers that build models to address questions in philosophy from papers that simply mention, discuss, or comment on models from a philosophical perspective. This distinction is important because philosophers of science can engage with the sciences without using any of the model-based methods that are common in many scientific disciplines. In such cases, philosophers do not contribute to "science in philosophy" in the sense of engaging with the sciences by drawing on model-based methods from scientific disciplines. Clearly, this is not to say that one way of engaging with the sciences is better than the other. But it is a distinction worth drawing, as the focus of this paper is not on philosophical work that mentions, discusses, or comments on models but instead on the use of model-based methods drawn from the sciences to tackle philosophical questions. According to our rubric, we therefore say that a paper builds a model when it uses a model to support a philosophical claim about the target of the model. In contrast, we say that a paper mentions, discusses, or comments on a model when it uses a model to support a philosophical claim about the model itself or its use. The distinction is thus akin to one that is often made in philosophy of language between mentioning a linguistic expression (cf. using a model to make a philosophical claim about the model or its use) and using the expression (cf. using a model to make a philosophical claim the model's target).

Consider, for example, Zollman (2007). In this paper, Zollman explicitly borrows model-based methods from economics to represent and study a community of scientists. Using computer simulations, Zollman finds that a community of scientists can be more reliable when scientists are less aware of their colleagues' experimental results and that there is a trade-off between the reliability and the speed with which the community reaches the right answer on a scientific question. This is paradigmatic case of a paper that uses model-based methods because it extends and adapts previous models to support a philosophical claim about the target of its model—namely, the behavior of a community of scientists. Similar examples include Huggett and Wüthrich (2018), Tanaka et al. (2020) and Weatherall et al. (2020): in all these cases, models are used to support philosophical claims about their targets.

In contrast, consider Bokulich (2003). Bokulich's focus in this paper is on quantum maps: models used to study the relationship between classical and quantum mechanics. She explores the use of these models by arguing that quantum maps belong to a family of "horizontal models": models that are built not from theory or experimental results, but from analogies with models in neighboring disciplines. This is a paradigmatic case of a paper that does *not* use model-based methods because it mentions, discusses, and comments on models to support a philosophical claim that is not about the target of model or group of models but rather about the use of such models in a scientific subdiscipline—in particular, the use of quantum maps in quantum chaos research. Similar examples include Weisberg (2007), Oreskes et al. (2010), Gelfert (2011), as well as other papers that invoke models to support philosophical claims about the models themselves or their use.

Although this rubric allows us to draw a distinction between using and mentioning models, it is also important to emphasize that this is of course not the only possible rubric. At the same time, not any rubric will do. A choice of rubric is a consequential methodological decision. But as it is often the case with such decisions, it is not one that can be made in the absence of a goal or purpose. Given the goal of isolating the use of model-based methods drawn from scientific disciplines to address philosophical questions, we have chosen to use one rubric among many that allows us to single out papers that are representative of the phenomenon we are interested in. But we acknowledge that other researchers might have decided to use a different rubric to study the same phenomenon. Comparing results obtained on the basis of different rubrics would in fact be a worthwhile project.

After manually labelling the entire set of 500 papers, we split labelled papers into two sets: a training set with 400 papers, and a testing set of 100 papers. This is common practice in classification tasks because it allows us to estimate parameters and the accuracy of the classifier using separate data sets. That is, the training set was used to estimate the parameters used in the classification task and thus to train the naive Bayes classifier; the testing set was used to determine the accuracy of the classifier. Since labels were manually assigned to all papers in both the training and the testing set, we could determine how often the classifier assigned the correct label to papers in the testing set given the parameters estimated using the training set. In this way, it was possible to estimate the accuracy of the classifier in the entire dataset by using the accuracy of the classifier in the testing set.

We then labeled the remaining $8,717$ papers with the help of the naive Bayes classifier; the entire dataset with labelled papers is available in the repository provided below. With papers thus labelled and sorted into communities, we were then able to track how the proportion of papers using model-based methods changed over time in each community. However, the mere presence of a significant difference in the proportion of papers using model-based methods does not tell us whether the observed change was due to random chance or a preference for a particular methodology. To determine whether and in what communities there has been a preference for the use of model-based methods, we therefore built a model to represent the cultural evolution of methods in each of the communities of papers that we identified within philosophy of science.

Models that represent the cultural evolution of epistemic practices in academic communities are now common in philosophy—for landmark papers and recent examples, see Weisberg and Muldoon (2009), Bruner (2013), Bright (2017), Zollman (2018), O'Connor (2019), and Heesen (2019). A central assumption of these models is that researchers choose what epistemic practices to pursue by copying others. These models therefore assume that epistemic communities change via a process of cultural evolution in which epistemic agents are the focal unit of analysis. Although this is a plausible assumption to make in many cases, in other cases it is also reasonable to suppose that cultural evolution takes place in a population of artefacts—for a discussion of these alternative formulations of cultural evolution, see Ramsey and De Block (2017). As our data pertains to papers and not researchers, we choose artefacts as our focal unit of analysis and thus assume that cultural evolution takes place in a population of research artefacts—i.e., papers.

To do so, we built a model for the cultural evolution of methods in philosophy of science using a modeling framework known as the Wright-Fisher model—for an early mathematical treatment and a recent philosophical discussion, see Wright (1931) and Clatterbuck (2015). Similar versions of the Wright-Fisher model have already been used to study the evolution of cultural artefacts, such as words (Sindi and Dale, 2016; Newberry et al., 2017; Karsdorp et al., 2020). In its simplest form, the Wright-Fisher model assumes that evolution takes place in a population with discrete types and discrete generations. In every generation, individuals are chosen to reproduce in proportion to how many individuals of each type there are in the population. Upon reproduction, all individuals die and a new generation is born.

In our case, the two discrete types correspond to the two types of papers (papers that use model-based methods and papers that do not) and discrete generations correspond to the publication year of research papers—see Figure 1. Every year, papers are chosen to reproduce in proportion to how many papers of each type were available in the previous year. The population of papers grows over time because papers never leave the population: for simplicity, we assume that are no retractions and thus that papers never leave the publication record once they have been published. This model for the cultural evolution of methods in philosophy of science therefore represents change over time in the methodological profile of the discipline under the assumption that the methods used in papers are chosen on the basis of what methods were used in papers published previously—for details on the model, see Appendix 4: Wright-Fisher Model.

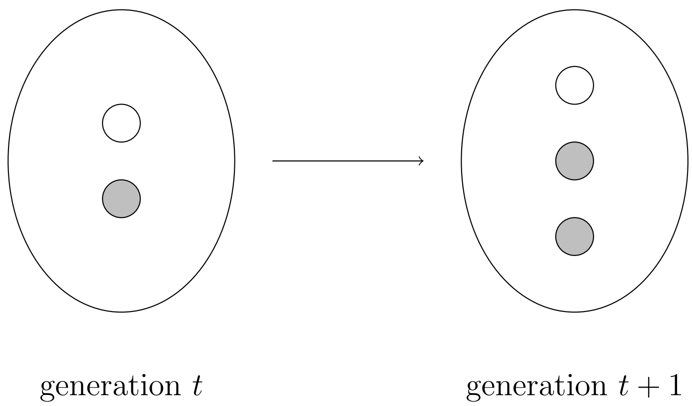

Figure 1: Example population in a model with two discrete types (white and grey), discrete generations, and growing population size. With selection given by the coefficient s, the probability that a population of size $N_t = 2$ and one grey individual transitions to a population of size $N_{t+1} = 3$ and two grey individuals is equal to $Pr(i_{t+1} = 2|i_t = 1) = 1 \cdot p^1 q^0$, with $p = \frac{1+s}{2+s}$ and $q = \frac{1}{2+s}$.

For all its simplicity, this model is useful because it allows us to estimate the strength of selection for or against the use of model-based methods within communities of papers in our bibliographic coupling network. To do so, we use the technique of maximum-likelihood estimation (Bolker, 2008). That is, we first calculated the probability of observing the actual trajectory of a community of papers given different values of s. We then took our estimate $\hat{s}$ to be the value of s that maximizes this probability—for details on how we used our cultural evolutionary model to estimate the strength of selection using maximum likelihood, see Appendix 5: Maximum-Likelihood Estimation.

To validate the results we obtained using the Wright-Fisher model, we followed Fletcher et al. (2021) and ran a regression analysis to determine whether the use of model-based methods has grown over time within each community. Publication year was the continuous independent variable and the dependent variable was whether a paper used model-based methods. This regression analysis provides a robustness check on our estimates of selection because it indicates whether there was a significant increase in the proportion of papers using model-based methods without the assumptions that go into the Wright-Fisher model. To determine whether there was an overall increase in the use of model-based methods, we also ran a regression analysis with the same independent and dependent variables for the entire dataset—that is, disregarding community membership.

Having described our data, the methods used to analyze it, and the model we use to represent our object of study, we present our results in the next section. Data sets and scripts are available anonymously at:

`https://osf.io/tm6v9/?view_only=2bb42691e5be4f9ca6ceec87b4860e48`

3 Results

Using Web of Science records for all papers written in English and published in the main philosophy of science journals between 2000 and 2020, we first built a bibliographic coupling network based on the cosine similarity between $tfidf$ scores for every pair of research paper matching our search criteria. This network contained $N = 9,217$ nodes corresponding to research papers and over one million edges between them. To simplify analysis, we therefore discarded edges with weight less than 0.05. The remaining network had the same number of nodes and $110,540$ edges.

This network had 390 connected components. In graph theory, a connected component is a set of nodes such that one could traverse from any node in the set to any other node in the same set via the edges connecting them. In informal terms, a connected component is thus a set of nodes that hang together and that is isolated from nodes outside the set. The largest connected component had $8,782$ nodes with $110,474$ edges between them. None of the 202 remaining components had more than eight nodes, with most of the components being singletons. To focus on papers that are representative of the discipline as a whole, we selected the largest connected component in the network; all other components were excluded from subsequent analyses.

By searching for a partition that maximizes network modularity, we then detected 20 distinct communities of papers in the largest connected component. Of these communities, four communities with fewer than 100 papers were excluded to ensure that enough data was available for community-level analysis. Overall, the remaining 16 communities contained $8,654$ papers (Table 2). Communities varied greatly in size (ranging from 171 to $1,162$ papers) and in number of edges (ranging from 424 to $17,637$). The mean number of papers per community was 541 ($s.d. = 265$), with a mean number of $5,2645$ edges ($s.d. = 4,141$).

To identify the main research topics in each community of papers, we extracted all keywords occurring in every paper in a given community and ranked them according to frequency of occurrence. Communities were labeled with the three most common keywords. We further identified the paper with the highest degree centrality in each community, degree centrality being the sum of the weights of all edges of a given node. We then assigned a topic to each community on the basis of most common keywords and most

Table 2: List of communities with assigned topic, most common keywords, most central paper, number of nodes (i.e., papers), and number of edges between papers.

No.	Topic	Keywords	Paper	Nodes	Edges
1	History	*Kant, Newton, Immanuel Kant*	Kochiras (2011)	171	424
2	Logic	*epistemic logic, belief revision, dynamic epistemic logic*	Renne (2008)	233	1,210
3	Mind	*perception, theory of mind, social cognition*	Kulvicki (2007)	322	1,184
4	Confirmation	*confirmation, probability, coherence*	Brössel (2015)	326	4,697
5	Teleology	*predictive processing, function, teleology*	Barrett (2014)	357	2,674
6	Social	*social epistemology, values in science, interdisciplinarity*	Biddle (2013)	383	3,192
7	Quantum	*quantum mechanics, Bohmian mechanics, entanglement*	Lewis (2007)	414	2,959
8	Evolution	*natural kinds, concepts, evolution*	Ramsey (2013)	421	3,712
9	Metaphysics	*grounding, ontology, vagueness*	Tugby (2021)	533	4,736
10	Models	*models, representation, representation*	Ducheyne (2012)	570	5,222
11	Relativity	*structural realism, general relativity, quantum mechanics*	Ainsworth (2011)	618	5,034
12	Decision	*decision theory, probability, rationality*	Shaw (2013)	645	7,329
13	Realism	*scientific realism, realism, incommensurability*	Doppelt (2005)	659	7,140
14	Knowledge	*knowledge, belief, epistemology*	Alspector-Kelly (2011)	892	11,382
15	Truth	*truth, semantics, propositions*	Bangu (2013)	948	5,696
16	Explanation	*explanation, causation, understanding*	Fagan (2012)	1,162	17,637

central paper. As shown in Table 2, the largest communities address questions in general philosophy of science, such as the nature of knowledge (No. 14, Knowledge: "*knowledge, belief, epistemology*"), truth (No. 15, Truth: "*truth, semantics, propositions*"), and explanation (No. 16, Explanation: "*explanation, causation, understanding*"). The smallest communities address topics in the history of philosophy (No. 1, History: "*Kant, Newton, Immanuel Kant*"), logic (No. 1, Logic: "*epistemic logic, belief revision, dynamic epistemic logic*"), the philosophy of mind (No. 3, Mind: "*perception, theory of mind, social cognition*").

These communities closely correspond to the topics that Malaterre et al. (2021) identify taking a topic-model approach. In particular, the communities on Mind, Confirmation, Social, Quantum, Evolution, Relativity, Knowledge, Truth, and Explanation seem to correspond to homonymous topics in Malaterre et al. (2021). At the same time, the community on History seems to correspond to the topic on Classics in Malaterre et al. (2019), whereas Logic seems to partly correspond to Formal and Language, Teleology to Neuroscience, Metaphysics to Philosophy and Property, Models to Explanation and Scientific Theory, Decision to Agent-Decision and Game-Theory, and Realism to Scientific Theory. Despite similarities between these two sets of communities, it is important to keep in mind that neither the data nor the methods used in both studies are the same. So differences in the number and composition of these communities should be expected.

Next, we classified each paper as to their methodology ("uses a model-based method" vs. "does not use a model-based method") using a multinomial naive Bayes classifier. Out of the $N = 9,217$ research papers in our sample, the classifier identified $1,215$ papers that use model-based methods. This represents 13.2% of all papers in the dataset. Despite its simplicity, the classifier performed quite well in the classification task. Its overall accuracy was 0.92, meaning that the classifier was able to correctly label 92% of papers in the testing set. The overall accuracy alone does not specify the rate of false positives (i.e., papers that were incorrectly tagged as using model-based methods) and the rate of false negatives (i.e., papers that were incorrectly tagged as *not* using model-based methods). Yet, a closer look at the classifier's error rates revealed that its false-negative rate was 0.23 and that its false-positive rate was 0.057. The classifier's overall performance was therefore quite high: despite a relatively high false-negative rate, the classifier behaved quite conservatively as it had a very low false-positive rate; results reported below therefore represent an underestimate of the role that model-based methods play in philosophy of science.

The resulting classification allowed us to determine the proportion of papers using model-based methods in each community. Some communities contained a very high concentration of papers using model-based methods, while other contained almost none. For example, a community on general topics in philosophy of science contained almost as many papers that use model-based methods as papers that do not (Figure 2, left; Decision: "*decision theory, probability, rationality*"). At the same time, one community of papers in the philosophy of physics contained a moderate amount of papers using model-based methods (Figure 2, center; Relativity: *"structural realism, general relativity, quantum mechanics"*). And a community of papers addressing questions about the metaphysics of science contain very few papers using model-based methods (Figure 2, right; Metaphysics: *"grounding, ontology, vagueness"*).

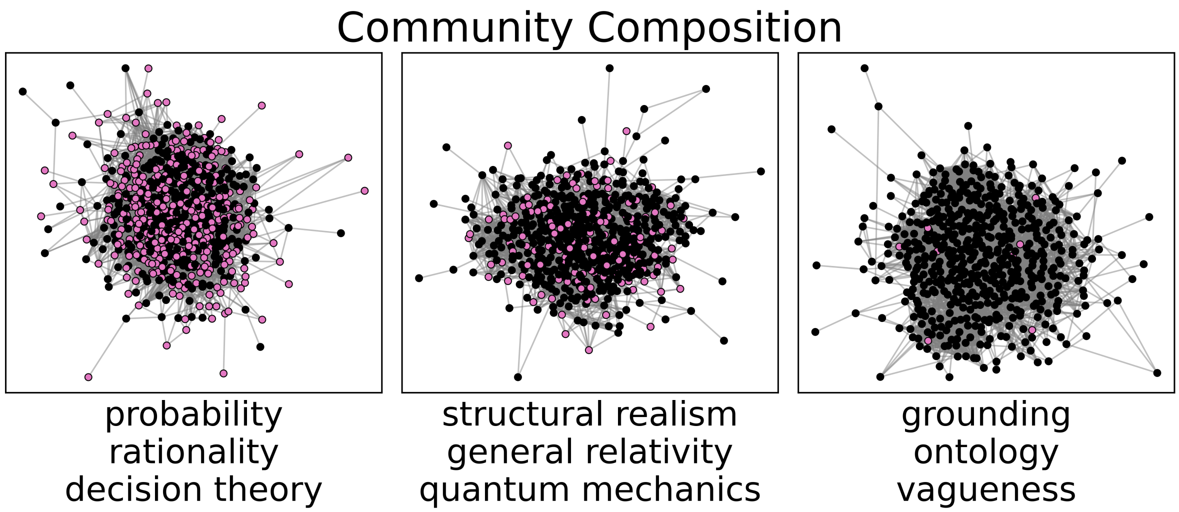

Figure 2: Examples of communities with low, moderate, and high concentration of papers using model-based methods. Nodes correspond to papers. Colors indicate papers that use model-based methods (*pink*) and papers that do not use model-based methods (*black*). For ease of visualization, all edges are shown regardless of their weight.

We also considered how the composition of each community changed during the time period analyzed. Again, we found variation across communities. Although the share of papers using model-based methods remained constant over the past two decades in some communities, it increased considerably in others (Figure 3). For example, there were very few papers using model-based papers published each year in some communities on the metaphysics of science (Figure 3, diamond; Metaphysics: *"grounding, ontology, vagueness"*). Other communities—for instance, in the philosophy of physics—had for the most part a constant number of papers using model-based methods published each year (Figure 3, cross; Realism: *"structural realism, general relativity, quantum mechanics"*). Still other communities experienced an increase in the number of papers using model-based methods over the time period, such as the community on decision theory (Figure 3, circle; Decision: *"probability, rationality, decision theory"*. However, it is not possible to determine looking at this change alone whether change was due to a general preference for such methods (i.e., cultural selection) or simply the result of chance fluctuations in the methodological profile of the community (i.e., random drift).

To answer this question, we built a model representing the cultural evolution of methods in philosophy of science from 2000 to 2020. By fitting the observed data to the model, we were able to determine the strength of selection for or against the use of model-based

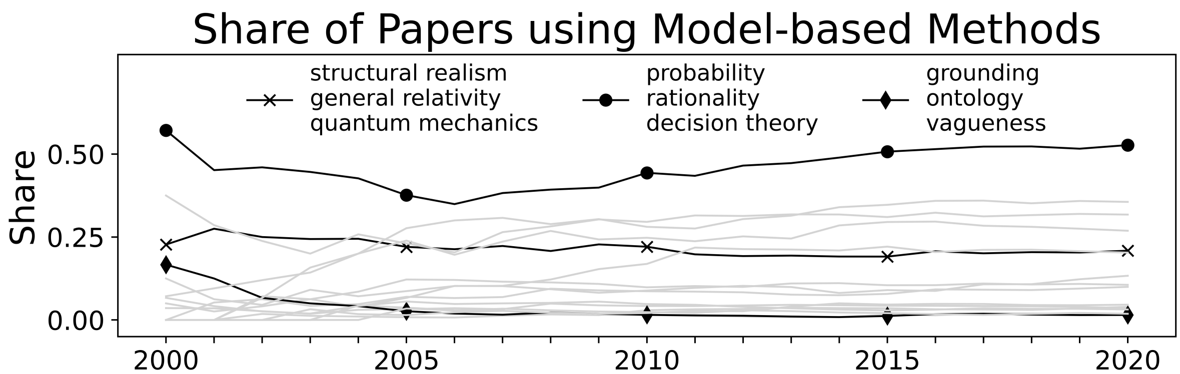

Figure 3: Time-series of the proportion of papers using model-based methods. Shown are examples communities on topics in the philosophy of physics (cross, *"structural realism, general relativity, quantum mechanics"*), general philosophy of science (circle, *"probability, rationality, decision theory"*), and the metaphysics of science (diamond, *"grounding, ontology, vagueness"*). All other communities are shown in gray.

methods. To do so, we used a technique for maximum-likelihood estimation: we determined the strength of selection by choosing the value of selection that maximizes the probability of the observed data. Using this technique, we inferred that selection favored papers using model-based methods in some, but not all communities.

In particular, we found three broad classes of communities (Figure 4). In the first class, communities have a substantial share of papers using model-based methods and selection for the use of such methods is high. This class encompasses communities such as those dealing with questions in decision theory (Decision: "*probability, rationality, and decision theory*") and the social dimension of science (Social: "*social epistemology, values in science, interdisciplinarity*"). A second class consists of communities in which there is again a significant share of papers using model-based methods but where absence of selection for the use of such models cannot be ruled out. Among these are communities addressing topics in the philosophy of biology (Evolution: "*natural kinds, concepts, evolution*") and logic (Logic: "*epistemic logic, belief revision, dynamic epistemic logic*"). And a third class consists of communities that did not experience strong selection for the use of model-based methods and that contain a very small share of papers using model-based methods. Examples include communities addressing issues in the metaphysics of science (Metaphysics: "*grounding, vagueness, and dispositions*") and the history of science (History: "*Kant, Newton, Immanuel Kant*"). Note that confidence intervals around selection estimates are wide in such communities, as the small share of papers using model-based methods makes it difficult to estimate the strength of selection accurately in these cases.

Note also that we were not able to detect selection against the use of model-based paper in any community. Although this merits further investigation, it is likely that this is at least in part due to limitations of our dataset: debates in the history and philosophy of

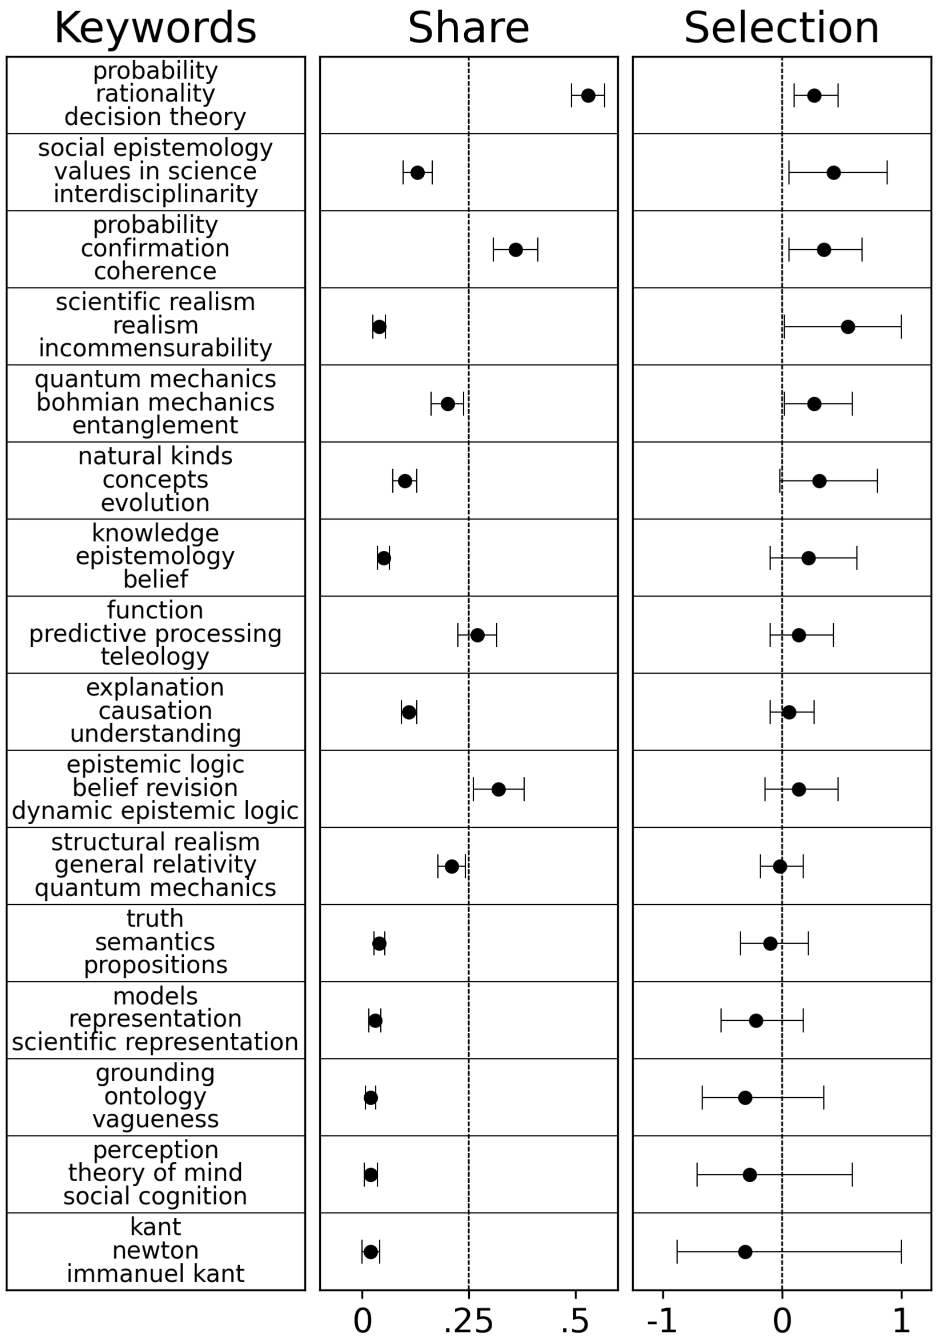

Figure 4: Most common keywords, proportion of papers using model-based methods, and maximum-likelihood estimate of selection. Error bars indicate two-tailed 95% confidence intervals. For estimates of selection $\hat{s}$, confidence intervals are given by values of s that satisfy the expression $\ell(s) - \ell(\hat{s}) \leq 1.92$, where $\ell(s)$ is the sum of log-likelihoods of the data given s.

science or in the metaphysics of science that tend to make less use of model-based methods are often published in journals that do not focus primarily on philosophy of science. For this reason, papers in our dataset may over-represent debates that experienced selection for the use of model-based methods during the time period considered. At the same time, it is important to emphasize that the model we used to estimate selection takes the population size of each community into account. This is clearly a virtue of this approach: large changes in the composition of the population are less likely to be due to selection in small populations than in large populations. As the population grows over time, our model is therefore sensitive to the fact that early changes are less likely to be due to selection than changes that take place later on.

Table 3: List of communities with effect of publication year on share of papers using model-based method (β) and associate p-value ('*' indicates p-value is less than 0.05); for comparison, estimate of s are also included ('*' indicates estimate for which the 95% confidence interval does not include zero.

	Topic	β	p	s
1	HISTORY	.011	.82	$-.31$
2	LOGIC	.016	.316	0.14
3	MIND	$-.016$	.61	$-.27$
4	CONFIRMATION	.03	.027*	.35*
5	TELEOLOGY	.011	.39	.14
6	SOCIAL	.042	.023*	.43*
7	QUANTUM	.013	.28	.27*
8	EVOLUTION	.024	.11	.31
9	METAPHYSICS	$-.01$	.72	-0.31
10	MODELS	$-.018$	.35	$-.22$
11	RELATIVITY	.006	.55	$-.02$
12	DECISION	.03	.0015*	.27*
13	REALISM	.04	.025*	.55*
14	KNOWLEDGE	.019	.25	.22
15	TRUTH	$-.01$	.44	$-.1$
16	EXPLANATION	.005	.59	.06

To validate results obtained on the basis of our cultural evolutionary model, we ran a logistic regression with publication year as the continuous independent variable and whether papers used model-based methods as the dependent variable (Table 3). Overall, there was no effect of publication year on the share of papers using model-based methods ($\beta = 0.007$, $p = 0.018$). At the community level, however, results varied substantively. In some communities, there was a significant effect of publication year on the share of papers using model-based methods. This was the case in the communities on social epistemology ($\beta = 0.042$, $p = 0.023$; SOCIAL: "*social epistemology, values in science, interdisciplinarity*"), scientific realism ($\beta = 0.04$, $p = 0.025$; REALISM: "*scientific realism, realism,* *incommensurability*"), decision theory ($\beta = 0.028$, $p = 0.0015$; Decision: "*probability, rationality, and decision theory*"), and confirmation ($\beta = 0.03$, $p = 0.027$; Confirmation: "*probability, confirmation, coherence*"). In all other communities, there was no significant effect of publication year on the share of papers using model-based methods.

Results from the regression analysis are largely consistent with results obtained on the basis of our cultural evolutionary model. Community in which we detected positive selection for the use of model-based methods were generally also communities in which there was an effect of publication year on the share of papers using model-based methods over the time period. The only exception to this was a community on the philosophy of quantum physics (Quantum: "*quantum mechanics, Bohmian mechanics, entanglement*"), in which selection for the use of model-based methods was detected although the logistic regression indicated that there was no effect of publication year on paper methodology. Similarly, communities in which the absence of selection could not be ruled out because confidence intervals around estimates of s included zero were also communities in which there was no effect of publication year on the share of papers using model-based methods. Thus, it is likely that there was some preference for the use of model-based methods in some but not all communities of papers in philosophy of science in the time period analyzed.

4 Discussion

In this paper, we built a bibliographic coupling network of all research papers written in English and published in the main philosophy of science journals from 2000 to 2020. Using an algorithm for community detection, we identified the most prominent communities in the network. We then classified papers with respect to their methodology using a supervised classifier. Results indicate that the share of papers using model-based methods did not increase overall but that it did increase in some though not all communities during the time period. Applying a model of cultural evolution to our data, we found evidence that the observed increase in the use of model-based methods can be attributed to cultural selection in some communities; these results were largely consistent with results from a logistic regression of paper methodology on publication year. Yet, in other communities we cannot rule out that changes in the use of model-based methods was simply due to cultural drift—i.e., random chance. Although our results go to show that there is variation in the strength of cultural drift and selection at the community-level, understanding what drives the trajectory of individual communities would require investigating the complex and multifarious tangle of factors that affect the natural history of each community—which is beyond the scope of this paper.

These results corroborate recent findings about the changing use of philosophical methods. Tracking the use of formal methods in a prominent journal, Fletcher et al. (2021) find that the use of probability theory significantly increased from the first to the second decade of the 20^{th} century. Similarly, Mizrahi and Dickinson (2020) find that the use of deductive arguments in JSTOR publications became less common during the early

2000s while the use of inductive and abductive arguments gained in popularity. Taken together, these studies therefore suggest that a shift in philosophical methodology took place around the same time period during which we also observe cultural selection for the use of model-based methods in many communities within philosophy of science.

It is also worth noting, however, that Fletcher et al. (2021) consider a single journal dedicated to general topics in philosophy and that Mizrahi and Dickinson (2020) consider all philosophy publications available in JSTOR. Both approaches are valuable. But neither take communities of papers addressing similar issues to be the units of analysis—as we do in this study. Moreover, their methods cannot determine whether the observed change in the use of methods was due to a cultural bias for any such method or simply random drift. Our study therefore expands on previous results by showing that a similar methodological shift can be attributed to a preference for the use of model-based methods in many debates within philosophy of science during the beginning of the 20^{th} century.

Our results also contribute to painting a fuller picture of the relationship between philosophy of science and the sciences. While Khelfaoui et al. (2021) and Pradeu et al. (2021) show that philosophy of science can contribute to the sciences when philosophers produce scientific knowledge with the help of philosophical methods, our results suggest that scientific disciplines can also contribute to philosophy. In particular, this can occur when philosophers borrow methods from the sciences—such as model-based method—to address philosophical questions. The relationship between philosophy of science and the sciences should therefore not be reduced to either one of these two complementary dimensions, as the evidence suggests that both "philosophy in science" and "science in philosophy" are constitutive of the relation between science and philosophy.

An additional strength of our approach vis-à-vis previous studies is that we were able to employ digital tools that allow for the analysis of large datasets. Fletcher et al. (2021), for example, note that a limitation of their approach is the focus on a single journal. To overcome this limitation, they suggest that future studies could "use more computational approaches" (p. 19). This is the approach we take here, analyzing changes in the methodological profile of an entire subdiscipline.

At the same time, our approach affords us greater resolution. By building a bibliographic coupling network, we were able to study the behavior of communities within the subdiscipline. Bibliographic coupling networks are built on the plausible assumption that papers with a similar citation pattern address similar topics. Communities in a bibliographic coupling network thus correspond to clusters of papers that address similar research questions, representing different areas of inquiry within a subdiscipline. As the variation in the share of papers using model-based methods and in the strength selection across communities suggests, such communities are indeed an important unit of analysis.

More generally, our results may have normative implications for graduate education in philosophy. Graduate students in philosophy are typically required to take few courses on methodology. When there are requirements in place, they often mandate courses in logic. As already noted by Fletcher et al. (2021), however, continuing use of formal methods other than logic raises questions about the appropriateness of such requirements. This is especially so if the goal of graduate programs is to prepare students to contribute to debates in philosophy of science that rely heavily on model-based methods. In such cases, philosophy departments would do well to train graduate students or at least guide them in how to acquire proficiency in such methods.

Finally, our study contributes to the integration of two emerging bodies of work in philosophy of science that have been isolated from one another. On the one hand, philosophy has recently seen a proliferation of models to represent the social dimension of epistemic communities—examples include already mentioned work by Weisberg and Muldoon (2009), Bruner (2013), Bright (2017), Zollman (2018), O'Connor (2019), Heesen (2019), and many others. On the other, recent work in philosophy has also turned to bibliometric data to study a variety of questions about scientific disciplines and their communities—for a few representative examples, see Byron (2007), Wray (2010), Machery and Cohen (2012), Overton (2013), and Weingart (2015). But the former body of work has for the most part not taken empirical evidence into account, whereas the latter often lacks a solid theoretical understanding of the phenomena it describes. Despite recent calls for integrating both approaches in the study of epistemic communities (Martini and Pinto, 2017; Thicke, 2020), little has been done to remedy the issue.

Yet, we show here that it is possible to integrate model-based and bibliometric approaches in the study of epistemic communities. By coupling an analysis of bibliometric data with a model for the cultural evolution of methods in philosophy of science, a major benefit of our approach is indeed that we were able to obtain a deeper understanding of the causes driving changes in the methodological profile within the subdiscipline.

5 Conclusion

Philosophers of science have recently turned to bibliometric data to answer a vast array of questions about science, philosophy, and the relation between the two. In many cases, the use of bibliometric data sheds new light on philosophical accounts of particular academic fields, such as the philosophy of biology, evolutionary behavioral science, or the history and philosophy of science (Byron, 2007; Wray, 2010; Weingart, 2015). More recently, philosophers of science have also relied on bibliometric data to investigate the relation between philosophy of science and the sciences in particular (Malaterre et al., 2019, 2020; Khelfaoui et al., 2021; Pradeu et al., 2021). We contribute to this body of work by showing here that philosophers of science not only participate in the production of scientific knowledge ("philosophy in science"), but also draw on model-based methods from the sciences to address philosophical questions ("science in philosophy").

There are some limitations to our approach, however. For one, we made several simplifying assumptions during data analysis and model construction. Bibliographic coupling networks assume that papers with a similar citation pattern address similar questions. The community-detection algorithm we used assumes that there are sharp boundaries between communities. The classifier we used assumes that the choice of labels is binary and that the features used in the classification task are probabilistically independent. And our model of cultural evolution assumes that generations are discrete and that there is no mutation. While we were able to justify these assumptions in the context of this study, it would be important to investigate the effect of relaxing these assumptions in future studies. In particular, it would be interesting to examine the use of model-based methods in philosophy of science by considering other types of networks (e.g., co-citation networks), fuzzy community-detection algorithm, and more sophisticated classification schemes that do not assume a binary choice of labels or probabilistic independence between features.

It is also important to emphasize that there are different sets of methods that philosophers of science can borrow from the sciences when addressing philosophical questions. We chose here to focus on model-based methods. But philosophers of science also make use of survey-based and experimental methods, as well as digital techniques and tools for bibliometric data analysis. Although model-based methods certainly make up an important set of methods that philosophers of science can and often do borrow from the sciences, it would be interesting to consider the use of other methods as well.

Relatedly, there may be a trade-off between depth and breadth of analysis in bibliometric studies—a trade-off similar to the one between precision and generality that Levins (1966) famously described in the field of theoretical biology. Here, we addressed the long-standing question in philosophy of science of how science relates to philosophy using techniques of "distant reading" (Moretti, 2000; Pence and Ramsey, 2018). Such a broad, big-data approach is clearly valuable, as it allows us to analyze large datasets. But it precludes us from closely engaging with individuals authors and papers, something that a narrow approach would be better suited for. It would thus be interesting to complement the present study by taking a narrow approach to the study of "science in philosophy".

Appendix 1: Bibliographic Coupling Network

In a bibliographic coupling network, nodes represent papers and edges represent the similarity between pairs of papers. To build a bibliographic coupling network, we first calculated the term frequency and the inverse-document frequency of references for each paper. The term frequency is given by:

$$tf(p_i, r_j) = \frac{f_j}{\sum_k^n f_k} \quad , \tag{1}$$

where f_j is the number of times that a reference r_j occurs in the reference section of paper p_i. Given that references are listed only once in academic papers, $f_j = 1$ if paper p_i cites reference r_j and 0 otherwise. Hence, $\sum_k^n f_k$ is the total number of references in the paper and n denotes the total number of references in the corpus. The term frequency ranges in the semi-open interval $(0, 1]$, being low when a paper cites a particular reference among many other references and high when it cites a reference among few others.

The inverse-document frequency is given by:

$$idf(r_j) = log\left(\frac{N}{M_j}\right) \quad , \tag{2}$$

where N is the total number of papers in the corpus and M_j is the number of papers in the corpus that cite reference r_j. The inverse-document frequency can take any real value, being low when many papers cite the reference and high when few papers cite it.

We then combined the term frequency and the inverse-document frequency to obtain the $tfidf(p_i)$ score for each paper. The $tfidf$ score is given by:

$$tfidf(p_i) = < tf(p_i, r_1) \cdot idf(r_1), ..., tf(p_i, r_n) \cdot idf(r_n) > \quad , \tag{3}$$

where $tf(p_i, r_j)$ and $idf(r_j)$ are defined as before. Notice that while the term frequency and the inverse-document frequency are scalar quantities, the $tfidf$ score is a vector.

Next, we measured the similarity between every pair of papers using the cosine similarity between their $tfidf$ scores. The cosine similarity is given by:

$$cos(p_i, p_j) = \frac{tfidf(p_i) \cdot tfidf(p_j)}{||tfidf(p_i)|| \cdot ||tfidf(p_j)||} \quad , \tag{4}$$

where $||tfidf(p_i)||$ is the so-called Euclidean norm of a vector and is given by $||\vec{a}|| = \sqrt{a_1^2 + ... + a_\ell^2}$ for a vector of length ℓ. It ranges in the unit interval, with 0 denoting complete dissimilarity and 1 denoting complete similarity.

Appendix 2: Community Detection

To detect communities of papers that engage similar research questions, we used a method that finds discrete communities in a network by maximizing network modularity. Given a partition of the network into communities $c_1, ..., c_m$, the modularity of a network is:

$$Q = \frac{1}{m} \sum_{i,j} \left(cos(p_i, p_j) - \frac{k_i k_j}{2m} \right) \cdot \delta(c_i, c_j) \tag{5}$$

where $k_i = \sum_j cos(p_i, p_j)$ is the sum of link weights for paper p_i and $m = \sum_{i,j} cos(p_i, p_j)$ is the sum of link weights for all papers in the network. The delta function $\delta(c_i, c_j)$ is equal to 1 if $c_i = c_j$, meaning that the community c_i of paper p_i is the same as the community c_j of paper p_j; $\delta(c_i, c_j)$ is zero otherwise.

Appendix 3: Naive Bayes Classifier

To label papers with respect to their methodology, we used a multinomial naive Bayes classifier. Naive Bayes classifiers assign an item to a class by maximizing the following expression:

$$Pr(q_i | w_1, ..., w_m) = \frac{Pr(w_1, ..., w_m | q_i) Pr(q_i)}{Pr(w_1, ..., w_m)} \quad , \tag{6}$$

where $Pr(q_i | w_1, ..., w_m)$ is the probability of the item belonging to class q_i given that the item has features $w_1, ..., w_m$, $Pr(q_i)$ is the unconditional probability of the class, and $Pr(w_1, ..., w_m)$ is the unconditional probability of the features. Items correspond to papers, classes correspond to the two types of methods that a paper might use (model-based method vs. no model-based method), and features correspond to the number of times that a word occurred in a paper's abstract and the number of times that a last name appears in a paper's reference section. These numbers are integers because words can appear any number of times in the abstract and last names can appear any number of times in the reference section.

Appendix 4: Wright-Fisher Model

To build a model for the cultural evolution of methods in philosophy of science, we assumed that papers are chosen to reproduce in proportion to how many papers of each type were available in the previous year. The probability that an individual of a given type—say, papers that use model-based methods—will be chosen to reproduce is given by:

$$p = \frac{i_t \cdot (1 + s)}{i_t \cdot (1 + s) + j_t} \quad , \tag{7}$$

where i_t is the number of papers of that type in generation t, $j_t = N_t - i_t$ is the number of individuals of the other type, and s is the selection coefficient measuring the strength of selection. Generations correspond to publication years. The parameter s is positive when selection favors the focal type, negative when selection favors the non-focal type, and zero when selection does not favor any type.

Conversely, the probability that a paper of the other type—papers that do not use model-based methods—will be chosen to reproduce is given by:

$$q = \frac{j_t}{i_t \cdot (1+s) + j_t} \quad , \tag{8}$$

where terms are defined as before.

Further, we assume that the population of papers grows over time because papers never leave the publication record. The probability that a population with i_t papers of a given type in generation t transitions to a population with i_{t+1} individuals of the same type in generation $t+1$ is thus given by:

$$Pr(i_{t+1}|i_t) = \binom{N_{t+1} - N_t}{i_{t+1} - i_t} \cdot p^{i_{t+1}-i_t} \cdot q^{j_{t+1}-j_t}, \tag{9}$$

where $\binom{N_{t+1}-N_t}{i_{t+1}-i_t}$ is the number of combinations we can obtain by choosing $i_{t+1} - i_t$ individuals of the focal type in a group of $N_{t+1} - N_t$ individuals, $p^{i_{t+1}-i_t}$ is the probability that $i_{t+1} - i_t$ individuals of the focal type will be chosen to enter the population, and $q^{j_{t+1}-j_t}$ is the probability that $j_{t+1} - j_t$ individuals of the non-focal type will be chosen to enter the population. Expression (9) therefore gives the probability that a population with i_t papers of a given type will transition to a population with i_{t+1} individuals of the same type by growing from size N_t to size N_{t+1}.

Appendix 5: Maximum-Likelihood Estimation

To estimate the strength of selection (s) for or against the use of model-based methods, we used the technique of maximum-likelihood estimation. That is, we take $\hat{s}$ be the value that maximizes the following expression:

$$\hat{s} = \underset{s\in[-1,1]}{\operatorname{argmax}} \sum_{t=2000}^{2020} log\left(Pr(i_{t+1}|i_t)\right) \quad , \tag{10}$$

where $\hat{s}$ is the maximum-likelihood estimate of selection for or against the use of model-based methods in a particular community, $Pr(i_{t+1}|i_t)$ is given by expression (9), and the sum is over the entire time period considered here—namely, from 2000 to 2020. Note that we take the log of $Pr(i_{t+1}|i_t)$ simply to facilitate computation, as values for $Pr(i_{t+1}|i_t)$ can be very small. Note also that equation (10) correspond to the estimate of selection for a particular community, so $\hat{s}$ must be estimated separately for each community of papers.

References

Ainsworth, P. (2011). Ontic structural realism and the principle of the identity of indiscernibles. *Erkenntnis*, 75(1):67–84.

Alspector-Kelly, M. (2011). Why safety doesn't save closure. *Synthese*, 183(2):127–142.

Bangu, S. (2013). Indispensability and explanation. *The British journal for the philosophy of science*, 64(2):255–277.

Barrett, J. A. (2014). Description and the problem of priors. *Erkenntnis*, 79(6):1343–1353.

Biddle, J. (2013). State of the field: Transient underdetermination and values in science. *Studies in History and Philosophy of Science Part A*, 44(1):124–133.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008.

Bokulich, A. (2003). Horizontal models: From bakers to cats. *Philosophy of Science*, 70(3):609–627.

Bolker, B. M. (2008). *Ecological models and data in R.* Princeton University Press.

Bright, L. K. (2017). On fraud. *Philosophical Studies*, 174(2):291–310.

Brössel, P. (2015). Keynes's coefficient of dependence revisited. *Erkenntnis*, 80(3):521–553.

Bruner, J. P. (2013). Policing epistemic communities. *Episteme*, 10(4):403.

Byron, J. M. (2007). Whence philosophy of biology? *The British journal for the philosophy of science*, 58(3):409–422.

Chandrasekar, P. and Qian, K. (2016). The impact of data preprocessing on the performance of a naive bayes classifier. In *2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC)*, volume 2, pages 618–619. IEEE.

Clatterbuck, H. (2015). Drift beyond Wright-Fisher. *Synthese*, 192(11):3487–3507.

Doppelt, G. (2005). Empirical success or explanatory success: What does current scientific realism need to explain? *Philosophy of Science*, 72(5):1076–1087.

Ducheyne, S. (2012). Scientific representations as limiting cases. *Erkenntnis*, 76(1):73–89.

Fagan, M. B. (2012). The joint account of mechanistic explanation. *Philosophy of Science*, 79(4):448–472.

Fletcher, S. C., Knobe, J., Wheeler, G., and Woodcock, B. A. (2021). Changing use of formal methods in philosophy: Late 2000s vs. late 2010s. *Synthese*.

Frigg, R., Nguyen, J., et al. (2020). *Modelling nature: An opinionated introduction to scientific representation*. Springer.

Gelfert, A. (2011). Mathematical formalisms in scientific practice: From denotation to model-based representation. *Studies in History and Philosophy of Science Part A*, 42(2):272–286.

Griffiths, P. E. and Stotz, K. (2008). Experimental philosophy of science. *Philosophy Compass*, 3(3):507–521.

Heesen, R. (2019). The credit incentive to be a maverick. *Studies in History and Philosophy of Science Part A*, 76:5–12.

Huggett, N. and Wüthrich, C. (2018). The (a) temporal emergence of spacetime. *Philosophy of Science*, 85(5):1190–1203.

Karsdorp, F., Manjavacas, E., Fonteyn, L., and Kestemont, M. (2020). Classifying evolutionary forces in language change using neural networks. *Evolutionary Human Sciences*, 2.

Kessler, M. M. (1963). Bibliographic coupling between scientific papers. *American documentation*, 14(1):10–25.

Khelfaoui, M., Gingras, Y., Lemoine, M., and Pradeu, T. (2021). The visibility of philosophy of science in the sciences, 1980–2018. *Synthese*, pages 1–31.

Knobe, J. (2007). Experimental philosophy. *Philosophy Compass*, 2(1):81–92.

Kochiras, H. (2011). Gravity's cause and substance counting: contextualizing the problems. *Studies in History and Philosophy of Science Part A*, 42(1):167–184.

Kulvicki, J. (2007). What is what it's like? introducing perceptual modes of presentation. *Synthese*, 156(2):205–229.

Leitgeb, H. (2013). Scientific philosophy, mathematical philosophy, and all that. *Metaphilosophy*, 44(3):267–275.

Levins, R. (1966). The strategy of model building in population biology. *American scientist*, 54(4):421–431.

Lewis, P. J. (2007). Empty waves in bohmian quantum mechanics. *The British journal for the philosophy of science*, 58(4):787–803.

Machery, E. (2016). Experimental philosophy of science. *A companion to experimental philosophy*, pages 475–490.

Machery, E. and Cohen, K. (2012). An evidence-based study of the evolutionary behavioral sciences. *The British Journal for the Philosophy of Science*, 63(1):177–226.

Malaterre, C., Chartier, J.-F., and Pulizzotto, D. (2019). What is this thing called philosophy of science? a computational topic-modeling perspective, 1934–2015. *HOPOS: The Journal of the International Society for the History of Philosophy of Science*, 9(2):215–249.

Malaterre, C., Lareau, F., Pulizzotto, D., and St-Onge, J. (2021). Eight journals over eight decades: a computational topic-modeling approach to contemporary philosophy of science. *Synthese*, 199(1):2883–2923.

Malaterre, C., Pulizzotto, D., and Lareau, F. (2020). Revisiting three decades of biology and philosophy: A computational topic-modeling perspective. *Biology & Philosophy*, 35(1):1–25.

Martini, C. and Pinto, M. F. (2017). Modeling the social organization of science. *European Journal for Philosophy of Science*, 7(2):221–238.

Mayo-Wilson, C. and Zollman, K. J. (2021). The computational philosophy: simulation as a core philosophical method. *Synthese*, pages 1–27.

McCallum, A., Nigam, K., et al. (1998). A comparison of event models for naive bayes text classification. *AAAI-98 workshop on learning for text categorization*, 752(1):41–48.

Mizrahi, M. and Dickinson, M. (2020). Argumentation in philosophical practice: An empirical study. In *Evidence, Persuasion & Diversity. Proceedings of OSSA 12, 3-6 June 2020*.

Moretti, F. (2000). Conjectures on world literature. *New left review*, 1:54.

Newberry, M. G., Ahern, C. A., Clark, R., and Plotkin, J. B. (2017). Detecting evolutionary forces in language change. *Nature*, 551(7679):223–226.

Noichl, M. (2021). Modeling the structure of recent philosophy. *Synthese*, 198(6):5089–5100.

O'Connor, C. (2019). The natural selection of conservative science. *Studies in History and Philosophy of Science Part A*, 76:24–29.

Oreskes, N., Stainforth, D. A., and Smith, L. A. (2010). Adaptation to global warming: do climate models tell us what we need to know? *Philosophy of Science*, 77(5):1012–1028.

Overton, J. A. (2013). "Explain" in scientific discourse. *Synthese*, 190(8):1383–1405.

Pence, C. H. and Ramsey, G. (2018). How to do digital philosophy of science. *Philosophy of Science*, 85(5):930–941.

Pradeu, T., Lemoine, M., Khelfaoui, M., and Gingras, Y. (2021). Philosophy in science: Can philosophers of science permeate through science and produce scientific knowledge? *The British Journal for the Philosophy of Science*.

Ramsey, G. (2013). Organisms, traits, and population subdivisions: Two arguments against the causal conception of fitness? *The British journal for the philosophy of science*, 64(3):589–608.

Ramsey, G. and De Block, A. (2017). Is cultural fitness hopelessly confused? *The British Journal for the Philosophy of Science*, 68(2):305–328.

Ramsey, G. and De Block, A. (2021). *The Dynamics of Science: Computational Frontiers in History and Philosophy of Science*. Pittsburgh University Press; Pittsburgh.

Renne, B. (2008). Public and private communication are different: Results on relative expressivity. *Synthese*, 165(2):225–245.

Shaw, J. R. (2013). De se belief and rational choice. *Synthese*, 190(3):491–508.

Sindi, S. S. and Dale, R. (2016). Culturomics as a data playground for tests of selection: Mathematical approaches to detecting selection in word use. *Journal of theoretical biology*, 405:140–149.

Sprenger, J. and Hartmann, S. (2019). *Bayesian philosophy of science*. Oxford University Press.

Suárez, M. (2008). *Fictions in science: Philosophical essays on modeling and idealization*. Routledge.

Tanaka, M. M., Godfrey-Smith, P., and Kerr, B. (2020). The dual landscape model of adaptation and niche construction. *Philosophy of Science*, 87(3):478–498.

Thicke, M. (2020). Evaluating formal models of science. *Journal for General Philosophy of Science*, 51(2):315–335.

Tugby, M. (2021). Grounding theories of powers. *Synthese*, 198(12):11187–11216.

Weatherall, J. O., O'Connor, C., and Bruner, J. P. (2020). How to beat science and influence people: policymakers and propaganda in epistemic networks. *The British Journal for the Philosophy of Science*, 71(4):1157–1186.

Weingart, S. B. (2015). Finding the history and philosophy of science. *Erkenntnis*, 80(1):201–213.

Weisberg, M. (2007). Three kinds of idealization. *The journal of Philosophy*, 104(12):639–659.

Weisberg, M. (2012). *Simulation and similarity: Using models to understand the world.* Oxford University Press.

Weisberg, M. and Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76(2):225–252.

Wheeler, G. (2013). Models, models, and models. *Metaphilosophy*, 44(3):293–300.

Wray, K. B. (2010). Philosophy of science: What are the key journals in the field? *Erkenntnis*, 72(3):423–430.

Wright, S. (1931). Evolution in mendelian populations. *Genetics*, 16(2):97.

Zollman, K. J. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5):574–587.

Zollman, K. J. (2018). The credit economy and the economic rationality of science. *The Journal of Philosophy*, 115(1):5–33.