

Philosophy of Science Meets Quantitative Studies of Science

International workshop, University of Turin, 27-29 May 2024

Sala Lauree Terracini (Palazzo Nuovo, via Sant'Ottavio 20, Turin)

Abstracts

Monday 27

Mapping science and scientific change

How to map the changing conceptual space of the sciences?

Yves Gingras (Université du Québec à Montréal)

While archival research can give access to a sub-region of the intellectual field at a given time, methods to visualize and map what was called the “republic of letters” in the 17th require methods that can give a access to the global structure and not the local one. I will give example based on scientific correspondences of the 17th to the 19th century as well as using publications covering the whole of 20th century to show how co-citation analysis can be used to shed lights on the evolving global structure of the sciences as well as of their specialties. I will also discuss the difference in the sociological meaning of co-citations and bibliographic coupling to insist on the importance of constructing indicators that are adequate to the interpretations they are given.

Visual Analytic Studies of Science

Chaomei Chen (Drexel University)

The landscape of science is constantly changing and uneven in terms of attention, uncertainty, consensus, interdisciplinarity, just to name a few. To understand such a complex and dynamic landscape of science, one would need to examine patterns and trends across a wide range of granularity levels, from the novelty of a specific concept to a broader context at an interdisciplinary level, to synthesize insights and evidence from diverse, controversial, and even contradictory arguments, and to connect theories of science and the reality. I will outline a visual analytic framework of science of science and demonstrate what types of insights this framework of thinking can bring to us effectively and systematically and how researchers can benefit from such insights. In particular, I will highlight a series of illustrative cases: How did a research field shift its focus to the next research questions? What happened next to a field of research once a decade-long search for diagnostic evidence was found? What serve as the building blocks for interdisciplinary research? How can we tell it was the work of a recurring mechanism?

The “styles project” meets contextualized semantic shift detection

Arno Simons (Technische Universität Berlin), Adrian Wüthrich (Technische Universität Berlin),
Michael Zichert (Technische Universität Berlin)

Our contribution links the theoretical concept of styles, from the philosophy of science, to contextualized semantic shift detection (CSSD), an emerging methodology at the intersection of quantitative science studies, linguistics, and computer science.

More than 40 years on, Hacking’s “styles project” has lost none of its original appeal. The project continues to spark fruitful discussions and further developments (Sciortino 2023; 2017; Winther 2012; Hacking 2012; 1994; Crombie 1995; Fleck 1935). In essence, the styles project claims that the ways in which scientists think and act are shaped by certain styles, each of which introduces new types of objects, evidence, modalities, and possibilities.

Styles include and are expressed through new “sentences...things that were quite literally never said before” (Hacking 1994: 42). This naturally links the styles project to the question of how scientists write. So far, the emergence and transformations of styles has mostly been studied qualitatively and by close reading of selected texts. While this can offer detailed insights and precise contextualization, such an approach is limited in scope since a human researcher can only digest so many sources in a reasonable amount of time. Only more recently, quantitative work has investigated the actual usage of epistemic concepts in science by measuring the (co-)occurrence of certain “epistemic markers” (EM), i.e. words like “explanation” and “mechanism” (Malaterre and Léonard 2023; Mizrahi 2022; Overton 2013). But then again, this work neither refers to the styles project nor takes into account the toolbox of CSSD (see Montanelli and Periti 2023 for an overview).

Against this background, the recent advent of neural language modeling has opened up new avenues for the empirical investigation of styles in large text corpora. Building on both the EM and CSSD literatures, we employ contextualized word embeddings from large language models to study and compare the semantic shift of epistemic concepts. The general idea is to represent the concrete contextualized occurrences of epistemic concepts as high-dimensional vectors (embeddings), whose geometrical distances (e.g. cosine or euclidean) can be measured and compared to detect similarities, differences, and shifts in the meaning of the words (or sentences) in question.

Using our own domain-adapted bidirectional encoder for yielding contextualized word embeddings, we study the shifting meanings of markers such as “observation”, “experiment”, and “simulation” in two adjacent fields: astrophysics and high energy physics. According to the philosophy of physics we should expect to find that the meanings of these markers should differ between the two fields, because they are linked to (slightly) different epistemic strategies (Karaca 2023; Ableson 2023; Jacquart 2022; Morrison 2015; Heidler 2017; Knorr-Cetina 1999). We explore how the nuances of the concepts under investigation have undergone shifts over a 30-year period, spanning from 1990 to 2020, in a corpus of more than 600K articles from the arXiv.org, an established preprint server.

We also critically discuss the applicability and usefulness of these new methods for the future of the styles project and the interaction between philosophy of science and quantitative science studies more broadly.

Research evaluation and its effects

Towards responsible next generation metrics

Paul Wouters (University of Leiden)

The field of evaluating academic activities is vast, complex, and highly dynamic, as are the roles of any data and indicators used to support these evaluations. The upcoming LERU paper on Next Generation Metrics that I will present, explores how universities can and should use currently available metrics and data to assess their research processes, in conjunction with qualitative expertise and information, and describes the imperatives for the design and use of next-generation metrics.

We aim to develop a common understanding of what responsible next generation metrics are and what they can mean for universities. This report focuses on the use and advancement of next-generation metrics for responsible research evaluation, encompassing open science, societal impact, and innovation.

The LERU paper formulates recommendations for all stakeholders in and around the academic community to foster the development and use of responsible next generation metrics. This report also provides an extensive list of valuable resources on next-generation metrics. We place strong emphasis on the imperative of collaboration, both among universities and between universities and funding agencies. By leveraging international expertise in bibliometrics and informatics, universities will enhance their resilience, reduce dependence on a limited number of companies that provide (seemingly) pre-packaged research intelligence, and gain a greater capacity to shape their metric policies in alignment with their own missions, rather than relying solely on standard metrics and data availability.

The CoARA initiative for reforming research assessment: Does it make sense?

Giovanni Abramo (Universitas Mercatorum)

Under the European Research Area policy agenda 2020-2024, the establishment of the CoARA (Coalition for Assessment Reform in Academia) in 2022 marked a significant milestone. Comprising predominantly research institutions and funders, CoARA's primary objective is to overhaul research assessment methodologies, with a specific emphasis on transitioning towards qualitative judgment. This discussion delves into the theoretical, practical, and applicative aspects of the CoARA initiative, examining its implications. The aim is to provide insights and reflections on the intricacies of research assessment, a multifaceted discipline often approached and conducted with superficiality, lacking the necessary preparation and attention. The goal is to ensure that decisions in this realm are grounded in scientific principles and empirical evidence, avoiding potentially hazardous ideological deviations.

My talk will focus on some reflections, among which: The first: Is there a necessity for such a reformation? The second: Is it conceivable to envision a standardized outcome of the reformation that disregards the purposes, the applicative context, and the resources of research assessment? The third: The reform is advocated and expected to be executed by those who have formulated the

very research assessments considered in need of reform. Are they sufficiently competent? The fourth: Is a large-scale research assessment via peer-review economically feasible?

From truth to Impact: How metrics is changing the framework for the evaluation of science

Mario Biagioli (UCLA)

As an independent variable, time has been structurally woven into the production of knowledge, but not its evaluation. There is nothing relevant that connects Copernicus *De Revolutionibus* and 1543 – the year the Book was published. Outside of priority or patenting disputes, time is still “alien” to judgments of truth or quality. Time may enable a new paradigm from grow and take over, but the paradigm itself is a snapshot, not a movie, of a worldview. That created a series of epistemic and ethical problems that led to the development of scientometrics and metrics-based evaluation frameworks. The simplest way to avoid judgment bias is to get rid of evaluation, moving toward some kind of “mechanical objectivity” that focuses on impact rather than on quality judgments. But that means that time becomes crucial because impact is impact in time. Though the work of De Solla price and Eugene Garfield do not appear to be inspired by philosophical consideration, except for a tendency to see knowledge as divisible in “molecules, and quanta, which they take as the object of citations, turning them into endorsements and recognition of specific knowledge claims, not worldviews. While Garfield was explicitly presenting scientometrics as a tool to guide librarians in the difficult task of best allocating their limited budgets (that was the origin of the ‘impact factors of journals: if a JIF is high, it may be worth subscribing), Garfield seemed more attracted to *longue-durée* patterns of scientific and disciplinary output. But there is a way in which the development and application of metrics marks a rather central shift in the philosophy of science, the best example being the renaissance of the science of science, made possible by more sophisticate models and vastly larger datasets. There are already specialized units of the new science of science at Northwestern, Stanford, and Northeastern University.

I believe the intellectual hinge of this follows from the replacement of “impact” for “quality”. Quality is judgment-based but virtually instantaneous. I read a new article, like it, and find it to bring up good question. Of course this is a judgment, with all the biases (actual and implicit) that go with that. I am not a defender of judgment but am concerned that the alleged solution – impact – is not essentially better. (This is a somewhat complicated argument that I’m happy to go into if there’s interest). But what matters here is that impact is not a judgment but an effect, and effects follow from causes. In sum, impact does not simply exist like a metal in a specific ore but takes time to accumulate or grow. In my view, this conceptual framework is essentially different from that of traditional philosophy of science. The intriguing question, or one of them is whether the new quantitative social sciences can make a contribution or revise old tenets of Philosophy of science. So far, the questions being tackled are not terribly challenging. James Evans and his group have looked at the correlation between team size and innovation (arguing that smaller teams are more creative). However interesting, I don’t find these statements profound. Still, I am impressed by the fact that they are now possible. Garfield and De Solla Price are toasting somewhere, probably arguing what should happen next.

Tuesday 28

Studying science: philosophy of science and scientometrics in dialogue

Metasciences: threading the past and imagining an interwoven future

Cassidy Sugimoto (Georgia Institute of Technology)

The talk will provide a cursory overview of five metasciences—philosophy, history, sociology, economics, and science of science—with a focus on the shared research agenda across each of these areas. It will discuss how key questions from philosophy of science—such as demarcation, explanation, and justification—continue to be foundational concerns as new forms of metasciences are derived. It will go into depth on the quantitative areas of science study and the degree to which they theoretically draw upon their sister disciplines. It will then provide an overview of a future agenda for the metasciences, including issues such as social justice and sustainability, and how these require an integrated approach. Finally, it will consider the type of infrastructure that would be required to realize this interwoven vision.

Studying Science and Studying Science: Metascience and the Philosophy of Science

Charles Pence (Université Catholique de Louvain)

I doubt that I am the only speaker who was attracted to this meeting, at least in part, because it promises to bring two groups into dialogue who should be speaking to one another, but who largely have not: scholars working on metascience or quantitative studies of science, and scholars working in the philosophy of science. In this talk, I want to consider both general philosophical perspectives and some examples from recent work in my research group to try to explore the differences in perspective that might have led to this kind of separation. In part, I'll claim, this relates to differing stances about the normative and social purposes to which the study of science could be put; implicated as well are a host of issues involving conceptions of “naturalism” and the scope of “empirical” or “practice-driven” philosophy of science. I hope, at least in outline, to sketch some ways in which resolutions to these questions might encourage more fruitful dialogue.

Measuring a moving target - Innovation studies in practice

Andrea Scharnhorst (KNAW)

This paper is inspired by an innovation studies programme organized by Loet Leydesdorff and the author herself in 2002 [1]. Departing from a holistic and system-theoretical view of academic knowledge production, the main idea was to bridge conceptual views on innovation and concrete empirical ways to ‘measure the knowledge base’.

In the last decades, academia has fundamentally changed, not mainly in its essence (the fact that we think), but rather in the practice (the way we organise our thinking). By taking an evolutionary perspective, with the ongoing growth of the science system, we observe more differentiation, both in

terms of disciplines (cognitive level) as well as institutions (organisational level) [2]. Nonetheless, as it happens in the evolution of any complex system (and we consider the science system as one) we also observe counteracting trends. For instance: the emergence of new intermediaries and new coordinating levels. One of them are research infrastructures. Their role has changed from being purely supportive (e.g. responding to new technologies needed to perform research activities), to increasingly becoming co-creators in the knowledge production process [3].

This paper investigates to which extent these new coordinating structures can be captured in quantitative studies. It does so, by using concrete examples: (a) the DARIAH European Research Infrastructure; (b) a research data archive; and (c) the ‘Open Science movement’ (including Research Data management). Looking at ‘measuring attempts’ around these examples, this paper touches upon the fundamental questions of why, what and how aspects of knowledge production can and should be measure, using the perspective of the ‘Measuring the Knowledge Base’ programme introduced at the start of the paper.

Qualitative Bayes: Procedure for applying Bayes’ rule to theory confirmation using estimates of likelihoods and concepts from information theory

Henry Small (SciTech Strategies, Inc.)

Application of philosophy of science methods, such as Bayesian statistics, to episodes in the history of science as well as contemporary science has been hampered by the lack of a way to estimate subjective probabilities for individual actors or groups. Such probabilities, for example, might reflect the confidence an individual has in some proposed theory or the individual’s degree of certainty that a theory can explain or predict a particular experimental finding. No doubt such probabilities will vary widely over individuals in a specialty depending on their awareness of and attitude toward the relevant evidence and theory. I propose assigning such probabilities using a simple scale of values to assess an actors’ subjective state from the historical record. The method will be illustrated using the controversy over continental drift from the early 20th century. Connections to citation context analysis and information theory will be discussed.

The Key Journals in Philosophy of Science: An Update

K. Brad Wray (Aarhus University)

In 2010, I published a paper with the aim of determining which journals are the most important journals in philosophy of science (see Wray 2010). The paper was an attempt to find an objective measure of the influence of various journals in the field of philosophy of science. The data I drew on were from three sources, published in 1996, 2000, and 2008. So the findings that I reported on there, represent what the key journals were in the late 1990s and early 2000s. I propose to reexamine this question in light of new data and thus update our understanding of what the key journals are in the philosophy of science. I will rely principally on a source published in 2018 (see Saatsi 2018).

First, I aim to determine whether the list of most influential journals has changed since my 2010 study. Fields change, of course, but it is worth seeing if philosophy of science has changed much

since 2000. Significantly, since 2010, there is a new journal in the philosophy of science, the European Journal for Philosophy of Science, which began publication in 2011. It will be interesting to see if this journal has had a measurable impact on the field. As it is the official journal of the European Philosophy of Science Association (EPSA), the journal stands to have a strong impact, containing papers from their biennial conferences.

Second, I want to extend the study to include a study of the impact of books and book publishers in the philosophy of science. My earlier study focused narrowly of journals. But books have a significant impact in philosophy of science, and the humanities in general (see Wray and Bornmann 2015). Consequently, a study of the influence of book publishers would be useful. Further, I aim to determine what proportion of citations in the data are to books rather than journal articles.

Third, I want to revisit a question that I examined briefly in my earlier paper: is there evidence of an integrated field of History and Philosophy of Science, HPS? Or, do the two field, history of science and philosophy of science, operate quite independently of each other? In the earlier study, I found no evidence of an integrated field (though, see Weingart 2015). My sense is that the field has not changed in this respect. Philosophers of science work quite independently of historians of science and seldom draw on their work in their own publications.

Fourth, I want to reflect on the methods employed in this and the earlier study. This study, and my earlier study, rely on relatively small data sets.

A dialogue between philosophy of science and computational studies of science illuminates the crisis of fundamental physics

Lucas Gautheron (University of Wuppertal and École Normale Supérieure)

Fundamental physics is in crisis [1]. The Large Hadron Collider (LHC) failed any evidence supporting theoretical attempts to unify fundamental forces and to combine quantum mechanics and gravity, and it seems increasingly unlikely that particle colliders, which have long been the main driver of progress in the field, will further promote these goals. In this paper, I combine concepts from the philosophy of science with quantitative analyses of scientific literature to show, first, how the crisis of particle physics challenges the unity of the field and, second, how scientists have adapted opportunistically to the challenges met by collider physics. I conclude that concepts from the philosophy of science can inform quantitative approaches, but also that, in turn, computational methods can generate philosophical insights.

I start with the disunity of physics, which has been the subject of much philosophical discussion. While reductionists endorsed the primacy of “fundamental” physics, others promoted more horizontal and pluralistic views granting each area of physics much more relative autonomy. In this context, Peter Galison observed that high-energy physics is divided into “subcultures” (theory, experiment, and instrumentation) that entail distinct scientific communities and languages. He argued that these subcultures communicate through local exchanges of knowledge, within “trading zones” [2], and contested that there existed full translations of higher level theories into lower level theories. Similarly, I distinguish two theoretical subcultures in high-energy physics, namely “phenomenology” and pure “theory”, by revealing their linguistic and social entrenchment with

quantitative methods [3]. I analyze the trading zone between “theory” and “phenomenology” by examining the concepts exchanged in 3.7M citations across these two categories of literature, between 2001 and 2019. This shows that phenomenological models of the physical events explored in particle colliders have become less able to sustain exchanges between high theory and phenomenology since the start of the LHC, and that these two cultures are “coming apart”, thus challenging a teleological view of the unity of science.

Secondly, I investigate how the field has adapted to this “crisis” by conducting a quantitative longitudinal analysis of a cohort of 2195 high-energy physicists between the years 2000 and 2019 [4]. Physicists’ research portfolios are measured with a topic model and a novel Bayesian model of change in scientists’ research interests is proposed. The model reveals a shift from electroweak physics towards dark matter, and a disintegration of “string theory” into “black holes” and holographic dualities. Using a mathematical framework called “Optimal Transport”, I show that the reallocation of research efforts is shaped by “learning costs”, thus enhancing the utility of the scientific capital disseminated among scientists. Optimal Transport was initially formulated by mathematician Gaspard Monge in 1781, as he sought efficient strategies for displacing piles of sand to specific locations. Since then, this branch of mathematics has found applications in economics, physics, computer science, and more. I show that it suggests new quantitative approaches to epistemic change, but also philosophical insights on the role of conservatism in science.

Wednesday 29

Interdisciplinarity and isolation in scientific progress

The scientometrics of interdisciplinarity. From concepts to measurement

Wolfgang Glänzel (KU Leuven)

Interdisciplinarity (IDR) has become a main characteristic of contemporary scientific research and poses severe challenges to scientometrics. IDR is strongly connected with specific patterns of scientific communication in terms of requirements, mechanisms, effects, and impact, which results in various manifestations. Here we sketch approaches and methods developed to meet the challenges arising from this complexity in an appropriate manner. We show that the scientometrics of IDR needs to follow four important steps, starting from conceptualisation, followed by the definition and choice of appropriate subject classification and granularity, before observations can be quantified and, finally, a meaningful measurement can be implemented. We stress the possibility of different solutions, which, however, can be based on the same or at least similar methodology. The steps within the outlined framework are illustrated by an example from philosophy. This presentation forms the counterpart to the study by Pei-Shan Chi on isolation of research in philosophy together with which IDR forms another side of the same coin.

The paper is based on recent literature and, most notably, on results from research conducted by the ECOO bibliometrics group at KU Leuven.

Duality Phenomena in Scholarly Communication: Exploring Isolation through Scientometric Analysis in the Field of Philosophy

Pei-Shan Chi (KU Leuven)

This presentation builds upon the previous discussion on the scientometrics of interdisciplinarity, delving into the complementary aspect of duality phenomena in science. It shifts the focus from interdisciplinarity to examine the possibility of quantification characteristic aspects and thus the scientometric measurement of isolation, thereby completing the puzzle of scholarly communication by exploring both integration and separation within the academic landscape. Beginning with an overview of the background and motivation underlying the study of isolation in philosophy, we proceed to introduce the measurement of scientometrics, emphasizing important caveats and limitations. Drawing upon empirical findings from Chi and Conix (2022), we provide illustrative examples of measuring isolation across various subfields and topics within philosophy. In the conclusion of the presentation, we juxtapose the findings from both the interdisciplinarity and isolation perspectives, offering insights into the practical application of scientometric methods and their implications for homogeneity and heterogeneity within scholarly research and thus shedding light on an example of a quantitative approach to the duality concept.

Understanding the societal relevance of humanities research through peer-review: a quantitative study

Stijn Conix (Université Catholique de Louvain)

There is a growing demand from funding agencies and societal stakeholders for academic research to be societally relevant. Simultaneously, criticism abounds regarding the perceived lack of societal relevance in much of scientific research, in particular in humanities disciplines. This creates a strong incentive to measure the societal relevance of research. Unfortunately, this is notoriously difficult when it concerns fundamental research like in the humanities (Bornmann 2013). Traditionally, peer review has been a common tool for estimating societal relevance, particularly in grant funding applications, where it frequently serves as a key evaluation criterion. However, the use of peer review in this context raises concerns due to its susceptibility to biases (Guthrie et al. 2019), low interrater reliability (Erosheva, Martinková, and Lee 2021), and uncertainties regarding researchers' ability to accurately assess societal impact .

To deepen our understanding of the societal relevance of humanities research and to assess the effectiveness of peer review, we hired 40 PhD students from diverse backgrounds in both a pilot and main study (link to OSF project blinded for peer review). Participants evaluated the societal relevance of 850 book and paper abstracts spanning five humanities fields—philosophy, religion, history, literature, and linguistics. Operationalizing 'societal relevance' using Kitcher's notion of an *ideal committee* (Kitcher 2011), raters ranked abstracts in sets of 5 according to their societal relevance, and determined whether these abstracts would be selected for funding by a Kitcherian ideal committee. Utilizing multi-level generative Bayesian Thurstonian models (Lee and Ke 2022; Li, Yi, and Liu 2022), we estimated the impact of rater strictness and chauvinism on assessments of societal relevance, and explored differences in societal relevance between the five fields. Additionally, two

members of the research team coded all abstracts for eight content characteristics (e.g., ethics, discrimination of minorities, science, health and wellbeing), and the associations between these characteristics and relevance scores were investigated.

This presentation aims to present the study's findings as well as discuss the application of Kitcher's notion of the ideal committee as a framework for operationalizing and measuring the societal relevance of research.

Power and social capital in science and academia

Social capital in academia: measuring individual researchers' collaboration preferences versus habits

Alesia Zuccala (University of Copenhagen)

Social capital is a concept rooted in the social sciences, and depending on which scholar is asked for a definition, there is much disagreement pertaining to its theoretical and empirical utility. As a theory, it is used by researchers to explain something, and as a philosophy, social capital may be seen as fundamental to an individual's survival, if not to how a society or community works. In academia many studies have relied on social network theory and social network analysis as an approach to measuring social capital. Social network analysis focuses on a boundary set of actors, for example, co-authors of research articles, and what the co-author network reveals, as a result of the bridges and bonds (i.e. edges) between the actors. However, SNA, can be problematic because it fails to attend to the "*qualitative and behavioural dimensions of social capital*" (Martín-Alcázar et al., 2019). In a new study, we have taken a different approach, by investigating researchers' perceptions of social capital. Our hypothesis was that what academics prefer in terms of building social capital may not be what they habitually do. A questionnaire-survey was piloted at one university with 1,092 academics from six different faculties ((i.e., Health and Medical Science, Science, Social Sciences, Humanities, Law, and Theology), and based on Portes & Sensenbrenner (1993) definition- i.e., "*social capital is the expectations for action within a collectivity that affect the goal-seeking behaviour of its members*" – our focus was on academics' goal-seeking behavior of publishing new research.

Who are the gatekeepers of economics? Geographic diversity, gender composition, and interlocking editorship of journal boards

Alberto Baccini (Università di Siena)

The composition of the editorial boards of economics journals is explored in terms of geographic affiliation, institutional affiliation, and gender. Results highlight that the academic publishing environment is primarily governed by men affiliated with elite universities in the United States. Comparison of networks generated by all scholars, editorial leaders, and non-editorial leaders reveals significant structural similarities and associations among clusters of journals. Links between pairs of journals tend to be redundant, and this can be interpreted in terms of social and intellectual homophily within each board, and between boards of journals belonging to the same

cluster. The high concentration of editorial power poses a serious risk to innovative research in economics.

Funding and scientific progress

Diversity in Evaluative Criteria Enhances Scientific Inquiry. An Epistemic Landscape Model

Matteo Michelini (Technical University of Eindhoven and Ruhr University Bochum), Javier Osorio (Universitat Autònoma de Madrid)

Scientists may judge the quality of the same scientific approach in rather different ways even if they belong to the same community (Schindler, 2021; Ward, 2022). For example, Maxim and Van der Sluijs (2018) have shown that scientists studying the toxicity of Bisphenol A employ rather different evaluative criteria when judging the quality of designing choices. And, yet, they still collaborate and potentially learn from each other.

Naturally, one may wonder what is the best strategy for a scientist in this situation: should one mainly engage with peers who share their evaluative criteria, or should one seek out a more varied community? More broadly, what kind of epistemic environment is most effective, one where evaluative criteria are uniformly shared, or one characterized by diverse criteria?

Answering to these questions is not an easy task. On the one hand, individuals with similar criteria might benefit more from shared learning, as the solutions they develop are more likely to be relevant for each other. On the other hand, diverse evaluative criteria may generate more innovative ideas. Yet, despite extensive research on collective problem-solving (Smaldino et al., 2022; Wu and O'Connor, 2023), none of the studies have directly addressed these specific questions. They instead focus on scenarios where individuals assess the problem uniformly and aim to find a solution collectively (Hong and Page, 2004; Reijula et al., 2023; Wu, 2023).

This paper addresses these questions using an agent-based model, building upon the NK framework widely applied in biology (Kauffman and Levin, 1987), economics (Ganco, 2017), and philosophy (Reijula et al., 2023; Wu, 2023). We extend this model to represent a group of problem solvers who, while exploring the same set of solutions, may evaluate them differently and, consequently, pursue different ones.

We show that, perhaps surprisingly, problem solvers with moderately diverse criteria are on average more successful than problem solvers in homogeneous groups. Heterogeneous communities cover a broader range of solutions, as an agent might find a solution that, though not beneficial for her, may be ideal for another. This diversity proves particularly beneficial when tackling complex problems.

Furthermore, we find that while homogeneous communities benefit from limited information flow, in line with existing results in the literature (Lazer and Friedman, 2007; Zollman, 2007, 2010), heterogeneous ones thrive from continuous information exchange. Agents with different evaluative criteria profit greatly from learning about the progress made by others.

This qualifies further the received view on the ideal communication structure of a scientific community (Frey and Seselja, 2020; Rosenstock et al., 2017; Zollman, 2007). Therefore, we argue

that more empirical work on the extent to which problem solvers actually use different criteria is needed to understand which communication structure is more beneficial (Wu and O'Connor, 2023). Additionally, our results contribute to the increasingly large literature on epistemic diversity (Fazelpour and Steel, 2022; Holman et al., 2018; Hong and Page, 2004; Wu and O'Connor, 2023), by drawing attention on an overlooked type of diversity, that is diversity of evaluative criteria.

It is crucial to recognize that this ABM is highly idealized. The abstract nature of the model allows us to simulate possible scenarios and to promote theory building and hypotheses exploration. However, to enhance its accuracy, calibration against empirical data would be valuable (Martini and Fernandez Pinto, 2017). This would fine-tune the parametric space to reflect real-world behaviors and outcomes more closely. In particular, we identify several potential sources of empirical data for this purpose, such as bibliometric data or corpus linguistic analyses.

Overcoming conservatism in funding decisions: the selection procedure as a promising area of intervention

Thomas Feliciani (Polytechnic University of Milan), Chiara Franzoni (Polytechnic University of Milan)

Funding decisions skew against innovative research – a phenomenon called conservatism in science funding (Lee et al., 2013; Wessely, 1998). Despite efforts to enable innovative research via dedicated funding calls, grant peer review often falls short in identifying the most innovative submissions (Luukkonen, 2012; Veugelers et al., 2022). Furthermore, conservatism disproportionately affects disadvantaged groups of applicants such as women and early-career scholars (Hoppe et al., 2019; Mancuso et al., 2023; Wang et al., 2018). Consequently, a paradox emerges whereby research articles from funded projects may be less innovative than those that do not report any funding (Shin et al., 2022).

The highly competitive nature of funding calls can explain some of the obstacles for innovative research. Review panels must often draw the funding line among proposals that seem equally worthy of funding (Bornmann et al., 2010; van den Besselaar et al., 2018). To make these difficult determinations review panel discussions often turn risk-averse. Instead of searching for reasons why a given project should be funded, panel discussions may do the opposite, searching for reasons to reject the proposal. The inherently risky nature of innovative projects can thus be seen as a weakness and be the reason for rejection (Lane et al., 2022). Similarly, panel discussions striving to tease apart equally-promising proposals may over-rely on bibliometric indicators. But while conventional research gets more citation in the short term, innovative research accumulates citations in the long term (Stephan et al., 2017). These factors – and panel discussions specifically – put highly innovative proposals at a disadvantage.

Our contribution explores and evaluates alternatives to traditional peer review panel discussions, seeking selection procedures that can curb conservatism. Two interventions are promising: (1) randomization via partial lotteries (Fang & Casadevall, 2016; Shaw, 2023), and (2) aggregation methods for algorithmically combining reviews into funding recommendations (Feliciani et al., 2022; Hastie & Kameda, 2005). Randomization and some types of aggregation methods, we argue, might

be robust to – and compensate for – conservative and risk-averse reviewer evaluations, ultimately leading to funding decisions that enable innovative projects.

Through counterfactual experiments we retrospectively assess whether alternative funding procedures would have led review panels to select more innovative projects. We use data from Novo Nordisk Foundation from 23 funding calls between 2015-2021 comprising full-text proposals, applicant information, reviewer evaluations and final funding decisions for 1533 proposals. For each call, we compare funding decisions made in panel discussions with counterfactual decisions generated from combinations of funding lotteries and bias-robust aggregation methods. For each real or counterfactual review panel we measure the average degree of novelty among funded proposals. For generalizability and external validity we run our counterfactual analyses under various conditions (e.g. by assuming different funding rates) and explore and compare different definitions and operationalizations of novelty (Arts et al., 2023; Shibayama et al., 2021).

By evaluating the effectiveness of partial lotteries and aggregation methods in curbing conservatism, our study promises actionable interventions to enhance innovation-focused funding calls. It also prompts further inquiry into the differential impact these interventions might have on different dimensions of innovation.