

Balancing Specialization and Adaptation in a Transforming Scientific Landscape

Lucas Gautheron^{1,2}

Abstract

How do scientists navigate between the need to capitalize on their prior knowledge through specialization, and the urge to adapt to evolving research opportunities? Drawing from diverse perspectives on adaptation, this paper proposes an unsupervised Bayesian approach motivated by Optimal Transport of the evolution of scientists’ research portfolios in response to transformations in their field. The model relies on 186,162 scientific abstracts and authorship data to evaluate the influence of intellectual, social, and institutional resources on scientists’ trajectories within a cohort of 2 108 high-energy physicists between 2000 and 2019. Using Inverse Optimal Transport, the reallocation of research efforts is shown to be shaped by learning costs, thus enhancing the utility of the scientific capital disseminated among scientists. Two dimensions of social capital, namely “diversity” and “power”, have opposite associations with the magnitude of change in scientists’ research interests: while “diversity” is associated with greater change and expansion of research portfolios, “power” is associated with more stable research agendas. Social capital plays a more crucial role in shifts between cognitively distant research areas. More generally, this work suggests new approaches for understanding, measuring and modeling collective adaptation using Optimal Transport.

*lucas.gautheron@gmail.com

ORCID: 0000-0002-3776-3373

¹Interdisciplinary Centre for Science and Technology Studies (IZWT), University of Wuppertal, Germany

²Département d’Études Cognitives, École Normale Supérieure, Paris, France

Keywords adaptation; specialization; science of science; cultural evolution; computational social science; optimal transport

1 Introduction

Scientists are subject to conflicting incentives. On the one hand, they must work within the realm of their expertise, where they can most effectively exploit their prior knowledge and compete with peers; this conservative preference for familiar research topics is at the root of specialization. On the other hand, scientists are simultaneously compelled to revise their research interests to engage with more promising research areas in order to benefit from more exposure or to secure funding. Thus, in some instances, specialization is at odds with the need to adapt to the decline of certain research opportunities and the growth of new ones. How do scientists navigate the trade-off between specialization (i.e. the concentration of their intellectual resources within a narrow cognitive range) and adaptation (i.e. the need to adjust these resources to new realities)? This conflict differs from the “essential tension” between “tradition” and “innovation” proposed by Kuhn [1], or that between “exploration” and “exploitation” [2], which have both been explored quantitatively in previous works [3–9]. First, “adaptation” is not tantamount to innovation or disruption, for it can be a conformist move (e.g. as a result of a bandwagon effect [10]). Moreover, unlike “exploration”, adaptation is not identical to a search strategy in a static landscape [11], but rather the convergence towards a new state more congruent with current realities. Disruptions due to breakthroughs in Machine Learning or challenges due to climate change urge to understand how scientists adapt to changing circumstances. Therefore, the present paper investigates scientists’ responses to changes in their field (whether driven by epistemic or institutional

factors), and the effect of their capital (intellectual, social, or institutional) on their ability to adapt. Drawing insights on adaptation from cultural evolution and institutional change, we develop an unsupervised Bayesian approach to analyze changes in scientists’ research agenda while measuring the effect of “capital” (intellectual or social) on their individual trajectories. The model is applied to a cohort of high-energy physicists between the years 2000 and 2019, a time during which the historical driver of progress in the field – particle accelerators – have been contested by emerging astrophysical experiments, thus transforming the landscape of opportunities.

Our approach reveals trends in the field: the boom of dark matter research – fueled by shifts away from the physics of neutrinos and the electroweak sector – and the partial disintegration of string theory into the study of black holes and holography/dualities. More importantly, this analysis also shows that changes in scientists’ research portfolios are shaped by learning costs, as scientific communities adapting to new circumstances address an “Optimal Transport” problem by reallocating research efforts efficiently. Optimal Transport is a mathematical framework initially concerned with the optimal displacement and allocation of resources [12–14], and has since then found wide-ranging applications. We show that it also provides a characterization of scientists’ behavior, as driven by the need to maximize the utility of their scientific capital under changing circumstances. Moreover, the comparative analysis shows that two dimensions of social capital, namely “diversity” and “power” [15], have opposite associations with change. While “diversity” of social capital – the extent to which scientists have access to diverse cognitive resources via their collaborators – is correlated with greater change and further diversification of scientists’ research interests, “power” – roughly speaking, the size of their network – is associated with more stability in their research interests. Social capital has a stronger association with transfers between research areas that are more cognitively “distant”. There is no discernible effect of institutional stability after controlling for academic age (although affiliation data is a bit noisy in the dataset). Overall, we contribute: i) a conceptual account of the features of change in scientists’ research interests; ii) a novel methodological approach that introduces a model of scientists’ trajectories connected to Optimal Transport and measures of intellectual capital, social capital, diversity, and power; and finally, iii) some empirical evidence from high-energy physics. More generally, this paper addresses the relative lack of empirical works within the body of literature that investigates science as a cultural evolutionary system [16]. It demonstrates that Optimal Transport provides an insightful description of certain aspects of collective adaptation, but also computational tools (such as Probabilistic Inverse Optimal Transport [17] and OT based measures of change) for measuring adaptive behavior, and more generally, mobility in physical and abstract spaces.

In what follows, Sections 1.1 and 1.2 summarize previous research and lay out the conceptual background on which the analysis rests, and Section 1.3 introduces the context of high-energy physics to which the model is applied. Section 2 elaborates the methodology: the data (2.1), the topic model approach for measuring authors’ research portfolios (2.2), the proposed measures of intellectual and social capital (2.3), and the model of scientists’ trajectories (2.4). Section 3 presents the results: i) the transfers of attention from one research area to another due to changes in the scientific landscape; ii) the structuring role of learning costs in the observed patterns of adaptation (3.1), and iii) the effect of capital (intellectual and social) and institutional stability on physicists’ strategies (3.2).

1.1 Empirical background

Several works have investigated the evolution of scientists’ research interests. For instance, by mapping the trajectories of 103,246 physicists over 26 years using the **American Physical Society (APS)** dataset and its topic classification (the **Physics and Astronomy Classification Scheme® (PACS)**), Aleta et al. [5] demonstrated that a majority of physicists gradually migrate to entirely different topics by the end of their careers while often staying within the same general area. They reveal differences between subfields of physics, such that “exploitation” (i.e. specialization, as opposed to the “exploration” of new topics) is especially prevalent in particle physics. Using the same data, Jia et al. [4] instead find an exponentially decaying distribution of changes in scientists’ interests. Previous works generally agree, however, on the graduality of change in research topics [4–6], as previously observed by Gieryn [18]. Recognizing that scientists typically investigate several research questions in parallel, Gieryn proposed four mechanisms of gradual change, including “accretion” (a problem is added to their “problem set”), “selective substitution” (one problem is replaced by another), and “selective disengagement” (one problem is neglected).

While [4–6] document the structure of changes in scientists’ interests, they do not relate these transformations to changes in epistemic and institutional context, or to the scientists’ incentives and resources. Tripodi et al. [7] have taken a step in this direction. Using the APS dataset, they show that physicists are more likely to explore areas to which they are connected via their collaborators, and highlight the crucial importance of collaborations in the expansion of research portfolios, especially for the exploration of research areas distant from one’s core specialization. However, their work does not primarily address the *transformations* of scientists’ research portfolios throughout time – they do not quantify “change” –, and they recognize the need for further longitudinal analyses. Finally, previous works have explored the connection between spatial mobility patterns and scientific mobility using gravity or radiation models [8, 9]. In particular, [8] used such methods to compare the characteristics of two types of scientists, “explorers” as opposed to “exploiters”.

The present paper complements previous works on changes in scientists’ research in several ways. First, since the focus is on adaptation strategies, the core of our approach is both comparative (as in [7], and unlike other previous works) *and* longitudinal (unlike [7], although their paper includes longitudinal robustness checks). Second, this contribution evaluates previously unexplored aspects, such as the choice between expansion or consolidation of research portfolios and the effect of affiliation stability. Third, this work relates the findings to the epistemic context of the field and its transformations by performing the analysis at a circumscribed scale (high-energy physics). Fourth, this work does *not* rely primarily on the APS dataset and PACS categories, on which most previous works depend [4, 5, 19, 7], or any other pre-existing classification of the literature. Research areas are clustered using an unsupervised topic model, such that this approach measures linguistic change, which is arguably a more direct proxy of cognitive change. Fifth, this paper is the first application of the Probabilistic Inverse Optimal Transport approach from [17], which provides an alternative to other approaches to mobility (e.g. gravity models). Finally, the proposed approach is grounded in theory, by operationalizing concepts such as capital [20, 21], and by exploiting theoretical insights from diverse approaches to “adaptation”.

1.2 Conceptual framework

A central dilemma of adaptation consists in choosing which resources to leverage among those already available (although those may be suboptimal or irrelevant under new circumstances) and which resources to abandon and replace with others (which may be inefficiently costly). By adapting gradually, scientists can strategically retain the benefits of “problem retention” (e.g., the exploitation of “accumulated skills and resources” in one area, or “of an established research network”, [18, p. 106]) while progressively investing resources in new research directions. This is illustrated in Figure 1, which represents the research portfolios of one scientist during two distinct time periods. Cells indicate the resources exploited by the scientist (e.g., concepts, models, methods, etc.) and colors indicate to what problem areas this knowledge is applied. Figure 1 shows how scientists can enter new research areas by repurposing certain resources to new ends [22, 23]. We call this strategy “*conversion*”, in reference to the typology of incremental institutional change proposed by Mahoney et al. [24]¹

Not all knowledge can be successfully applied to new research areas: as illustrated in Figure 1, entering new research areas typically requires “*layering*”², that is, the introduction of new concepts, models, or methods, on top of prior knowledge. The acquisition of knowledge entails learning costs, which can be partially avoided by collaborating with experts in the target domain [7]. Another mode of change is *displacement*, when the replacement of one research area for another involves significant neglect of prior knowledge. This may not be the preferred strategy, since it fails to take advantage of accumulated resources. However, certain knowledge may not apply to a new context, or sometimes there might be reasons to suspend a line of research in order to focus on more promising topics. Overall, we expect that these transformations will manifest themselves as changes in scientists’ linguistic behavior, i.e., as changes in the vocabulary of their publications. Generally, we expect an important amount of continuity in linguistic behavior, given the need to minimize cognitive learning costs by capitalizing on prior knowledge.

Scientists manage two kinds of assets when navigating the trade-off between specialization and adaptation:

¹Indeed, as shown in previous works on the transformations of high-energy physics facilities to photon science instruments [25–27], historical institutionalism can account for gradual adaptations with large cumulative effects taking place in response to scientific and technological change [28]. In this paper, we apply the typology of change to *individuals* rather than organizations.

²Again, borrowing the terminology from historical institutionalism.

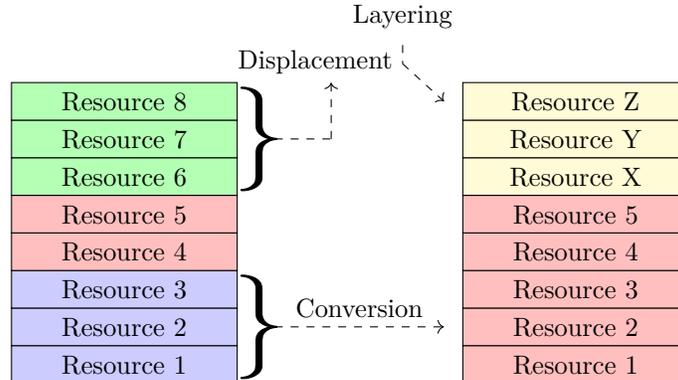


Figure 1: **Changes in a scientist’s “research portfolio” over time.** Colors designate research areas. Resources entail any intellectual or methodological assets that a scientist uses to investigate problems in each area. “Conversion” repurposes knowledge to new goals. “Displacement” is the replacement of certain research interests with little or no transfer of prior knowledge, as illustrated by the green area. Layering is the introduction of new research interests via the addition of new knowledge.

their own prior expertise, and the expertise to which they have access through their social network. Both constitute “capital” [21], i.e. assets that individuals accumulate and leverage in the competitive context of their field. “Capital” (whether “economic”, “cultural”, “social” or even “symbolic”, cf. [21]) defines the scope of scientists’ opportunities and therefore their ability to adapt. This paper considers the *intellectual capital* possessed by scientists in the form of scientific knowledge, and *social capital*. Measures that represent these concepts will be proposed, and their effect on the magnitude of transfers of attention across research areas will be evaluated. Emphasis will be put on the divide underlined in [15] between two dimensions of social capital, namely “power” (roughly speaking, network size in the present paper) and “diversity” (of cognitive resources). Group diversity is generally recognized as a factor of adaptation in an evolving environment or in the context of collective problem-solving [29–32]. “Power” is also plausibly associated with higher abilities.

While capital defines scientists’ opportunities, it is not sufficient to explain *why* scientists *do* turn to new research areas or not, which also requires understanding actors’ incentives and why they must respond to these incentives. Consequently, the present paper also considers the effect of institutional stability and academic age on migrations between research areas, since the need to respond to changes in the epistemic and institutional environment is presumably different for, say, tenured physicists versus postdocs. Moreover, we may assume that younger and older generations play different roles in cultural change and collective adaptation in general [33]. Finally, the effect of productivity will also be considered.

1.3 The case of high-energy physics: navigating a changing epistemic landscape

High-energy physics is a prime example for investigating adaptation in a transforming scientific landscape. As this field relied on the input of increasingly large particle colliders to achieve progress, it has accumulated considerable capital directed towards collider physics in the form of large infrastructure and complex knowledge. These efforts culminated in 2010 with the start of the **Large Hadron Collider (LHC)**, the largest accelerator ever. However, the LHC has found no evidence for anything that was not already predicted by the Standard Model of particle physics, and it is increasingly plausible that no future accelerator could ever find any evidence for new particles, leading to a situation of “crisis” [34]. Although the **LHC** will continue to take data for years and plans for successors are being discussed [35], some have speculated that particle physics as we know it has come to an end [36, 37], “the proscenium [being] captured by astrophysics and cosmology” instead [37]. However, according to physicist Mikhail Shifman the “pause in accelerator programs we are witnessing now is not necessarily [...] the end of explorations at [high energies]”; instead, such explorations “will continue, perhaps in a new form, with novel devices” [38]. Indeed, new experimental opportunities have emerged in parallel, including gravitational waves astronomy (since 2015 [39]), searches for dark matter of astrophysical origin in underground facilities, and more precise observations of the cosmic

microwave background (see Figure 10, Appendix A.1). Moreover, astrophysics seems to be increasingly replacing particle colliders in citations across experimental and theoretical high-energy physics [40]. The present paper investigates how high-energy physicists have adapted in reaction to these transformations.

2 Methods

2.1 Data

Our source is the Inspire HEP database [41]. It aggregates High-Energy Physics (HEP) literature from various sources, including the main scientific publishers and arXiv, and has been used in a few works [40, 42–45]. For the literature on HEP, it is more comprehensive than the often used APS dataset which is limited to a few journals³. Moreover, it implements both automatized and manual measures for the disambiguation of author names⁴, thus allowing careers’ analyses [44] (nevertheless, occasional misidentifications remain possible). The database also contains data on experiments; consequently, the evolution of the landscape of experimental opportunities can be retrieved (see Figure 10, Appendix A.1).

The analysis includes all papers from the categories “Theory-HEP” and “Phenomenology-HEP” (inspired from arXiv’s categories “hep-th” and “hep-ph”), to which most HEP publications belong, which amounts to $D = 186,162$ articles between 2000 and 2019. The minority of purely experimental high-energy physics publications are excluded: such papers are typically authored by thousands of collaborators, and authorship data provide no information about individual experimentalists’ specialization. Therefore, this paper documents how theorists and phenomenologists have adapted to the changes outline above.

For the longitudinal analysis, two time periods are considered. An initial phase (2000-2009) is used to infer a reference “research agenda” for each physicist in the cohort, as well as their intellectual and social capital. A late phase (2015-2019) is used to measure how each physicist’s research agenda has shifted in comparison to the initial time period, in the context of the changes outlined above. The five-year gap between these two periods allows to measure the cumulative effect of the transformations in the scientific landscape that have unfolded gradually between 2010 and 2015 (had they been sudden, we would not have introduced such a wide gap – see Figure 10, Appendix A.1), together with the effect of the capital accumulated *prior* to these transformations. Only physicists with ≥ 5 publications during each time period (2000 to 2009, and 2015 to 2019) are included, resulting in a cohort of $N = 2\,108$ physicists. This study therefore considers physicists that have remained dedicated to high-energy physics, thus revealing adaptation and “survival” strategies *within* HEP, excluding authors that exited the field. This author inclusion rule excludes scientists who publish very irregularly; however, although scientists who continuously publish are a minority, they make up most of the publications in their field⁵. We do not seek “representativeness”, but rather to achieve enough variance to uncover patterns among this particular cohort of productive high-energy physicists. The median academic age of the cohort was 23 years in 2015. 49% of these physicists have had an affiliation that spanned the entire time period (see Appendix A.2, Figure 11).

2.2 Measuring research portfolios

Research portfolios are evaluated in terms of the distribution of keywords (n-grams) that belong to each research area within the scientists’ publications’ abstracts. Instead of relying on PACS categories and citation data as in most previous works, research areas are extracted with a topic model that recovers latent “topics” within the corpus and their vocabulary distributions, while directly classifying keywords into separate research areas. Arguably, texts provide a more direct access to the kinds of knowledge leveraged by a scientist in their publications: linguistic flexibility implies cognitive flexibility and a low commitment to a specific body of knowledge. Moreover, this approach is applicable to a wider range of situations for which only textual data (even short texts) are available. For instance, PACS codes are only sparsely available in

³<https://journals.aps.org/datasets>

⁴Besides the use of “advance algorithms” of author-disambiguation, Inspire invites scientists to correct their own publication record on the website (https://twiki.cern.ch/twiki/pub/Inspire/WebHome/INSPIRE_background.pdf, June 2014)

⁵Less than 1% of scientists active in the years 1996 to 2011 have published every year during this period, and yet they are responsible for 47% of the publications [46]; the 13% of physicists with ≥ 16 publications between 1985 and 2009 account for 82% of publications [9].

the recent **HEP** literature. A coarse-grained classification of the literature into $K_0 = 20$ broad “topics” is performed. The number of topics is always somewhat arbitrary, but that topic models give some control over the cognitive scale is a feature, rather than a bug⁶: as Gieryn [18] puts it, “in such analyses [of problem change], empirical findings will in part reflect the defined scope of problem areas”, which is itself arbitrary. In our case, we would ideally like our clustering to be just fine-grained enough to measure the impact of the shifts in the landscape of experimental opportunities that we are interested in (the start of the **LHC**, the rise of new probes of dark matter and black holes, etc.). In this respect, $K_0 = 20$ turns out to be just sufficient to discern the effects of the transformations in **HEP** discussed in Section 1.3 as well as the observed evolution in the popularity of various kinds of experiments shown in Appendix A.1, Figure 10. Additional models were trained for robustness assessment, setting different values for K_0 (15, 20, 25). More coarse-grained models (using lower values of K_0) are typically less able to observe fine-grained patterns of adaptation to changing experimental opportunities, and the initial $K_0 = 20$ model is better at distinguishing black hole phenomenology from cosmological phenomenology than the most fine-grained model (see Appendix A.3.5, Figure 14).

We use an embedding model [47], a recent and straightforward approach that relies on pretrained embeddings representations of the n-grams and provides more reliable classifications for heavy-tailed vocabulary distributions than previous models such as Latent Dirichlet Allocation [48]. Given the coarseness of the clustering, Language Models were not deemed necessary. The model is trained on $D = 186,162$ abstracts published between 2000 and 2019. Tokens are extracted from the papers’ titles and abstracts by filtering n-grams between one and three words matching syntactic expressions susceptible of carrying scientific information (by designating concepts, models, methods, etc.), following the procedure from [40, 49]. Embeddings are learned using a skip-gram model in $L = 50$ dimensions (few are needed given the small size of the vocabulary, $V = 4,751$; nevertheless, some analyses are re-iterated with $L = 150$; see Appendix A.3, Figure 12) [50]. We obtain the topics listed in Appendix A.3, Table 1. Four of the 20 resulting topics regroup keywords that do not clearly refer to any specific research area (e.g. “paper”, “approach”) and correlate poorly with the **PACS** categories, which suggests their lack of scientific dimension. Consequently, the present analysis only considers the $K = 15$ remaining topics that designate actual research areas (as confirmed by their strong tendency to preferably cite themselves, cf. Appendix A.3.4, and their correlation patterns with the **PACS** classification, cf. Appendix A.3.6). In order to enhance the robustness of the topic removal process, we made sure all retained topics had a maximal loading on the **PACS** classification higher than that of all removed topics.

Then, we derive n_{dk} , the amount of keywords in the abstract of d that refer to a research area k (using the method described in Appendix A.3.2), and consequently $X_{a,k}$, the amount of times keywords (“resources”, i.e. concepts, models, methods, etc.) in relation to research area k have occurred in papers (co-)authored by a in the initial time-period (2000 to 2009). Mathematically, $X_{a,k} = \sum_{d \in [2000,2009], a \in A_d} n_{d,k}$, where A_d is the set of authors of a publication d . The matrix $Y_{a,k}$ is derived similarly, using publications from the later time period (2015 to 2019). Research portfolios are then normalized into distributions $x_{ak} \equiv X_{ak} / \sum_{k'} X_{ak'}$ and $y_{ak} \equiv Y_{ak} / \sum_{k'} Y_{ak'}$, thus encoding how scientists divided their attention during each period. This approach ensures that research portfolios are evaluated based on the frequency of keywords that belong to each “topic”, according to the idea illustrated in Figure 1. Therefore, this approach captures variations in the prevalence of different kinds of vocabulary (and thus bodies of knowledge) exploited in scientists’ publications. For purposes of illustration, Figure 2 shows the research portfolios of one physicist who migrated from neutrinos to dark matter physics, and of one physicist who maintained their research agenda over the time periods considered.

⁶Previous works based on the **PACS** categories have leveraged the different levels of this hierarchical classification system to investigate different scales.

⁷Physicist A’s personal website reads: “I am working on particle astrophysics and cosmology. In particular, *I am interested in dark matter problem in the Universe, and how to probe it using annihilation products such as energetic gamma rays and neutrinos.* [...] I started my research career by studying supernova neutrinos from various aspects [our emphases].” (<https://staff.fnwi.uva.nl/s.ando/eng/Research.html>). This is therefore an instance of “conversion” of prior knowledge to new purposes, one of the forms of change drawn from historical institutionalism represented in Figure 1.

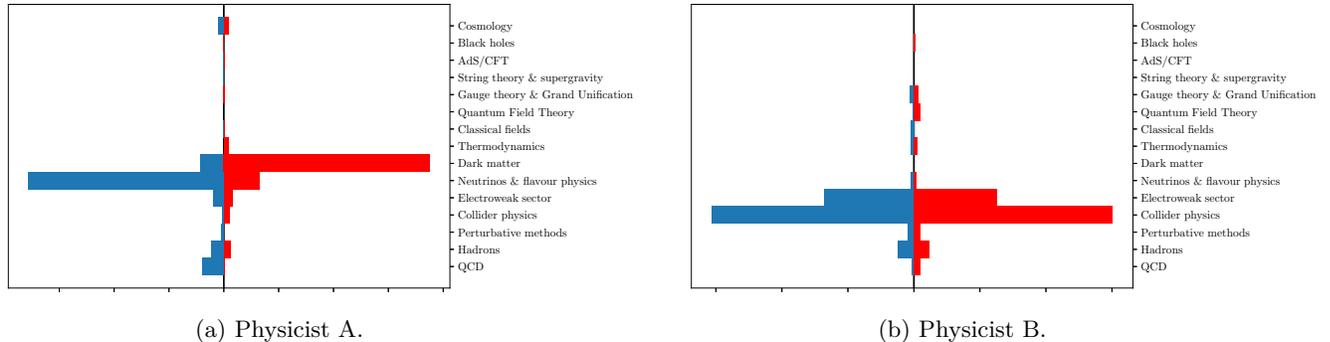


Figure 2: **Initial (in blue) and late (in red) research portfolios of two physicists.** Physicist A shifted their focus from neutrinos physics to dark matter⁷. Physicist B, on the other hand, pursued a remarkably stable research agenda.

2.3 Measuring capital

As shown by Schirone [51] in an extensive review of references to Bourdieu in bibliometrics, most mentions of capital focus on symbolic and social capital. Only a dozen works considered cultural/intellectual capital, and none of those proposed a measure that adequately captured the distribution of capital across epistemic domains. Therefore, an alternative unified approach for measuring the distribution of intellectual and social capital is proposed below.

2.3.1 Intellectual capital

Intellectual capital is represented by a vector $\mathbf{I}_a = (I_{ak})$ that measures the concentration of a 's intellectual resources in each domain $k \in \{1, \dots, K\}$. It is constructed in a way similar to \mathbf{x}_a , summing the contribution of keywords dedicated to each research area in the publications of each author between 2000 and 2009 (thus excluding publications that belong to the outcome research portfolio \mathbf{y}_a), except that publications are now weighed differently depending on the amount of authors. Indeed, publications with fewer co-authors convey more information about each author's own expertise. The weight is $\frac{1}{|A_d|}$, where $|A_d|$ is the amount of authors of publication d ⁸:

$$I_{ak} \propto \sum_{d \in [2000, 2009], a \in A_d} \frac{n_{d,k}}{|A_d|} \quad (1)$$

\mathbf{I}_a is normalized, such that $\sum_k I_{ak} = 1$; therefore, \mathbf{I}_a only captures the ways scientists divide their cognitive resources between research areas, rather than the “absolute magnitude” of their knowledge of each area (by contrast, the measure of semantic capital proposed in [52] measures total knowledge but cannot capture diversity).

2.3.2 Social capital

Many measures of scientists' social capital have been proposed [15, 51], the simplest being the amount of collaborators of a scientist (i.e. degree centrality in the co-authorship network [52]). Other measures revolve around *betweenness* centrality, which captures the extent to which an actor “bridges” a network (e.g. “brokerage”, i.e., the ability of an individual to overcome “structural holes” in a social network [53]). Abbasi et al. [15] distinguish two general approaches to social capital, depending on whether the emphasis is placed on “power” versus “diversity”. Measures of social capital (as those discussed in [15]) typically represent social capital by single scalars; however, social capital has multiple dimensions. In fact, according to Bourdieu [20], “the volume of social capital possessed by a particular agent [...] depends on the extent of the network of links that he can effectively mobilize, and on the volume of capital (economic, cultural or symbolic) possessed

⁸A justification for this weight is that the probability that a given author has been responsible for introducing any particular concept or method present in the paper is $\mathcal{O}(1/|A_d|)$.

by each of those to whom he is linked”. In that respect, social capital can come in different forms depending on the resources being leveraged via one’s network. In the following, we focus on the intellectual dimension of social capital, which we represent by a vector \mathbf{S}_a defined as the sum of the intellectual capital of a ’s collaborators, weighted by the strength of their relationship:

$$\mathbf{S}_a \equiv \sum_{c \in \text{co-authors}(a)} w_{ac} \mathbf{I}_{c \setminus a} \quad (2)$$

$\mathbf{I}_{c \setminus a}$ is the intellectual capital of c , evaluated by excluding papers co-authored with a (in order to disentangle the effect of an author’s own knowledge and that available to them via their collaborators). Collaborators outside the cohort are taken into account. The weight w_{ac} , which represents the strength of the relationship between a and c , is defined as:

$$w_{ac} \equiv \max_{d|\{a,c\} \subset A_d} \frac{1}{|A_d| - 1} \quad (3)$$

Where A_d is the set of the co-authors of a paper d . This weighing scheme – inspired from [54] – captures the fact that a paper with, say, two co-authors, signals a stronger relationship between the authors than a publication with a dozen authors.⁹ However, it does not take into account the recency and frequency of collaborations.

2.3.3 Diversity and power

Measures of “diversity” (and “power”) can be readily derived from \mathbf{I}_a (and \mathbf{S}_a). A common measure of diversity is the Shannon entropy H [55]. Let $D(\mathbf{I}_a) = \exp H(\mathbf{I}_a)$ be the diversity of intellectual capital (and $D(\mathbf{S}_a) = \exp H(\mathbf{S}_a)$ that of social capital). Roughly speaking, these are measures of how many research areas scientists have divided their cognitive or social resources among. In the cohort, individuals typically have cognitive resources in several research areas ($\mu_{D(\mathbf{I}_a)} = 5.6, \sigma = 2.0$), and social capital is even more diverse ($\mu_{D(\mathbf{S}_a)} = 7.8, \sigma = 2.1$). In fact, intuitively, scientific collaborations enable individuals to take advantage of their group’s diversity. Furthermore, $D(\mathbf{I}_a)$ and $D(\mathbf{S}_a)$ are highly correlated ($R = 0.75$); indeed, individuals with more diverse expertise are more able to engage with diverse collaborators. Since the diversity of social capital is mostly expected to enhance individuals’ abilities when it exceeds that of their own knowledge, from now on, only *excess* social capital diversity $D^*(\mathbf{S}_a)$ (defined as the residuals of the linear regression of $D(\mathbf{S}_a)$ against $D(\mathbf{I}_a)$, by ordinary least squares) is considered¹⁰.

The “power” dimension of social capital is evaluated as the *magnitude* of social capital:

$$P(\mathbf{S}_a) \equiv \sum_k S_{ak} = \sum_{c \in \text{co-authors}(a)} w_{ac} \quad (4)$$

“Power” is therefore the amount of collaborators weighed by the strength of each relationship. Our measures of diversity and power depart from [15], which conflates diversity with network size and power with performance. By combining semantic and authorship data, our approach assesses diversity more directly.

Alternative measures of diversity and power are considered for robustness assessment. The alternative measure of diversity is based on Stirling’s index, and the alternative measure of power uses the notion of brokerage. All these measures are defined and compared in Appendix A.4.

2.4 Modelling trajectories

The model for the late research portfolio \mathbf{Y}_a is schematically illustrated in Figures 3a,3b,3c, and a more formal representation is given in Figure 3d. It captures the idea expressed in Figure 1 that research portfolios are transformed via strategic transfers of knowledge and attention from one research area to another. Occurrences of keywords that belong to each research area k in papers by a in the late time period, $\mathbf{Y}_a \in \mathbb{N}^K$, are assumed

⁹Assuming that in a collaboration, each author interacts with a constant amount of co-authors in practice (regardless of the total amount of co-authors), then the probability that they had interactions with one specific co-author in particular is $\propto \frac{1}{|A_d| - 1}$.

¹⁰This approach aims to address the difficulty of determining the direction of the causal relationship between social resources and research interests raised by Tripodi et al. [7].

to be drawn from a hierarchical multinomial logistic model. \mathbf{Y}_a results from a linear combination of the initial research portfolio, \mathbf{x}_a , and a mixing matrix θ_a that measures the fraction of attention redistributed from each research area to another. θ_a is drawn from a hierarchical process, thus capturing the “average” cohort behavior. Formally speaking, \mathbf{Y}_a is assumed to derive from a multinomial process involving linear combinations of (x_{ak}) :

$$\mathbf{Y}_a \sim \text{multinomial}\left(\sum_{k=1}^K x_{ak}\theta_{ak1}, \dots, \sum_{k=1}^K x_{ak}\theta_{akK}\right) \quad (5)$$

Where $\theta_{akk'}$ is the fraction of attention to a topic k by a that has been redirected to a topic k' . θ is a function of intellectual capital \mathbf{I}_a and social capital \mathbf{S}_a according to the following generalized linear model:

$$\theta_{\mathbf{a}k} = \text{softmax}(\beta_{ak1} + \gamma_{k1}I_{a1} + \delta_{k1}S_{a1}, \dots, \beta_{akK} + \gamma_{kK}I_{aK} + \delta_{kK}S_{aK}) \quad (6)$$

$\delta_{kk'}$ is the effect of the scientists’ social capital in a research area k on the magnitude of transfers from k to k' . Similarly, $\gamma_{kk'}$ is the effect of having more expertise in k' (intellectual capital) on shifts from k to k' . High values of the diagonal elements of γ would imply that physicists are more conservative towards research areas in which they concentrate more expertise. The coefficients $\beta_{akk'}$ encode the average behavior of the cohort plus individual deviations to the average behavior that are unexplained by the covariates^{11,12}. The ability of the model to predict individual trajectories is assessed in Appendix A.5, for various temporal segmentations of the initial and late research portfolios and topic models. The model is better at predicting individual trajectories for larger collective adaptive responses and longer time-scales. Indeed, this hierarchical model relies on scientists’ collective behavior to predict individual change; moreover, research portfolios evaluated over shorter time-scales are less informative and more prone to noise.

Our model is strongly connected to Optimal Transport [56, 57], which seeks optimal ways to “transport” an input distribution (say, $\mathbf{x} = (x_k)$) to a target distribution (e.g., $\mathbf{y} = (y_{k'})$) through “transfers” ($\theta_{kk'}$) across their components while minimizing a cost function $\sum_{k,k'} x_k \theta_{kk'} C_{kk'}$ (where $C_{kk'}$ is a cost matrix) [14]. The difference is that the proposed Bayesian approach estimates transfers $\theta_{kk'}$ by minimizing a likelihood rather than a cost function. However, the connection with Optimal Transport suggests an economic interpretation of the reallocation of research efforts, which will be leveraged in Section 3.1 to show that patterns of change in research interests are shaped by learning costs.

3 Results

Figure 4a shows the aggregate transfers of attention at the level of the cohort revealed by the model. The most obvious feature is the remarkable stability of most research areas: indeed, physicists’ conservatism toward their research area due to specialization is known to be particularly high in HEP [5]; late research portfolios are largely constrained by prior research interests: they exhibit *path dependence* [11] (using Inverse Optimal Transport, Section 3.1 shows that these patterns are structured by learning costs). Conservatism seems especially prevalent in the case of “collider physics”, a research area dedicated to knowledge specific to particle accelerators. Nevertheless, “dark matter” has doubled, fueled by a shift away from “neutrinos

¹¹The priors for this hierarchical model are:

$$\beta_{akk'} \sim \mathcal{N}(\mu_{kk'}, \sigma_{kk'}) \text{ for } 1 \leq k' \leq K - 1 \text{ and } \beta_{akK} = \mu_{kK} \quad (7)$$

$$\mu_{kk'} \sim \mathcal{N}(\lambda \times \nu_{kk'}, 1) \text{ (average behavior)} \quad (8)$$

$$\delta_{kk'} \sim \mathcal{N}(\delta_0 + \lambda' \times \nu_{kk'}, 1) \text{ (effect of social capital)} \quad (9)$$

$$\gamma_{kk'} \sim \mathcal{N}(0, 1) \text{ (effect of intellectual capital)} \quad (10)$$

$$\sigma \sim \text{Exponential}(1) \quad (11)$$

Where $\nu_{kk'}$ is the fraction of physicists with expertise in k (that is, with more intellectual capital than average in k) who also have expertise in k' . Priors must be thought thoroughly, as certain invariances can lead to identification issues – for instance, shifting μ by a constant does not change the likelihood.

¹²The fit is performed with Stan’s Hamiltonian Monte-Carlo sampler (HMC is better behaved than Gibbs for such problems).

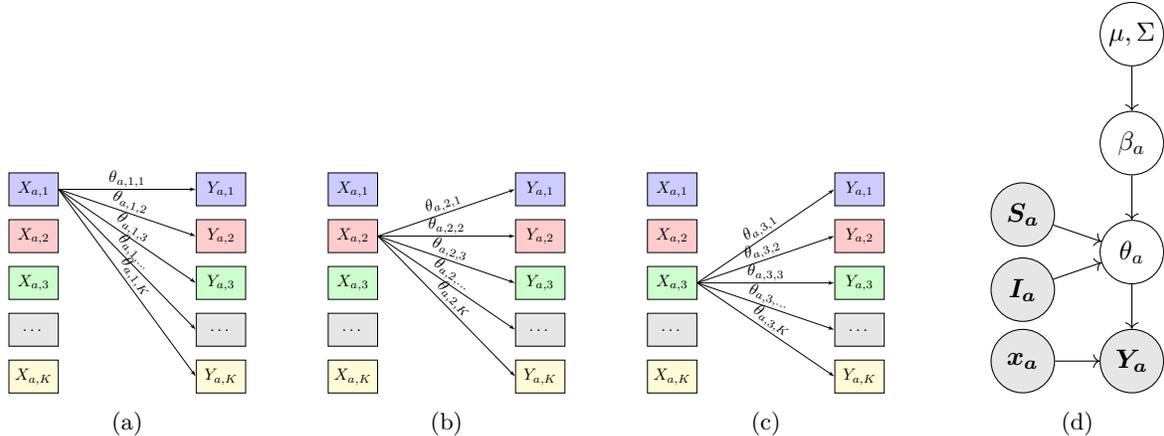


Figure 3: **3a,3b,3c: Transfers of attention across research areas.** x_a and y_a are the distributions scientist a 's attention across research areas in two consecutive time periods. $\theta_{\mathbf{a}k} = (\theta_{\mathbf{a}kk'})$ represents the fraction of the attention devoted by an author a to a topic k redirected to topics $k' \in \{1, \dots, K\}$, as scientists repurpose, expand or concentrate their knowledge. By definition, $\sum_{k'} \theta_{\mathbf{a}kk'} = 1$.

Figure 3d: Hierarchical model. θ_a is drawn from a hierarchical process, with intellectual and social capital (I_a and S_a) as covariates. Observed variables are represented in gray, latent variables in white.

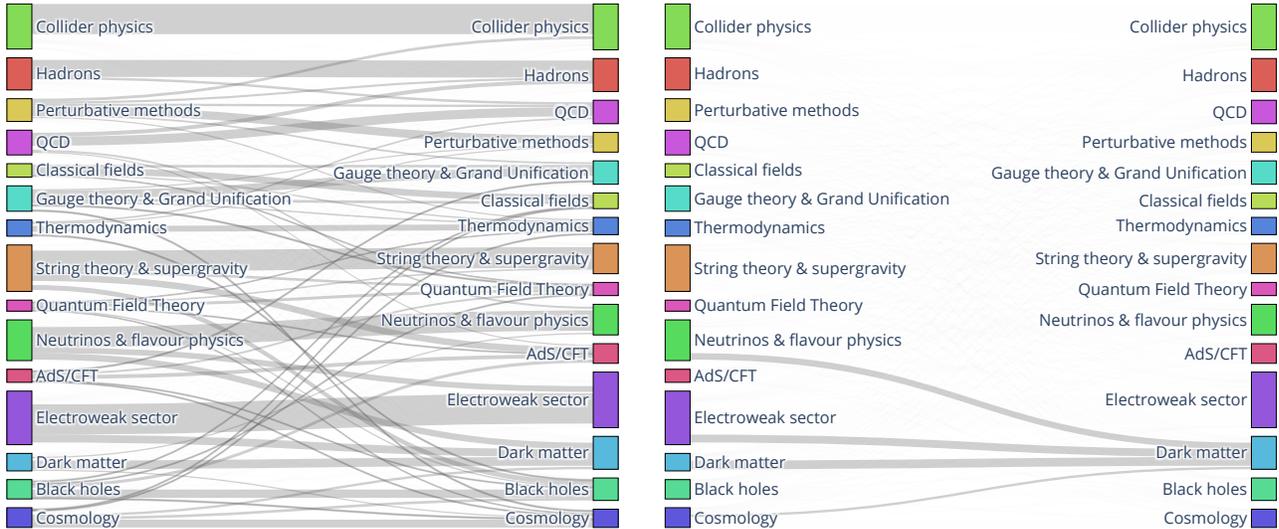
and flavor physics”, and “electroweak sector”, a phenomenal domain studied at the LHC (Figure 4b)¹³¹⁴. This confirms that the cohort has responded to changes in the landscape of experimental opportunities. Moreover, “string theory and supergravity” has declined in favor of “AdS/CFT” (a research program that explores dualities between theories of quantum gravity and certain types of theories of quantum field theory) and “black holes”¹⁵. Of course, this rough description must be considered with caution, as interpreting clusters from topic models is notoriously hard, and these topics in particular can regroup quite heterogeneous research programs.

While the paper focuses on two time periods (2000-2009 and 2015-2019), multiple alternative temporal segmentations can be considered. Figure 5 shows the transfers of research attention of a cohort of physicists across four time-periods of five years each. It reveals that the changes outlined above have unfolded rather gradually. The average diversity of physicists’ research portfolio ($\mu(D)$, the average of the exponentiated entropy of \mathbf{x}) during each five-year time-bin is also shown. It has gradually increased over the years ($P < 10^{-4}$): on average, physicists have *expanded* their portfolio. Interestingly, the average linguistic diversity of each individual paper increased as well (with a confidence level $P < 10^{-4}$) from 2.81 topics per paper (2000-2004) to 2.96 (2015-2019). This means physicists diversified their research portfolios in part by diversifying the knowledge leveraged *within* each of their individual papers (rather than solely by writing multiple papers on separate issues).

¹³The electroweak notably includes Higgs physics, which are very prominent at the LHC, where the Higgs boson was discovered.

¹⁴The migration of many particle physicists towards dark matter provides an explanation for the persisting schism between two research programs in fundamental physics, namely dark matter particle research and modified gravity. Both research programs seek to explain a shared set of anomalies in astronomical observations, and yet their communities communicate very little [58, 59]. Our approach suggests that particle physicists’ interest in dark matter is in great part motivated by the fact this is a natural extension of their previous research; particle physicists would therefore not consider the alternative to dark matter (modified gravity), given this topic that would make little use of their expertise.

¹⁵This converges with physicist Peter Woit’s controversial assessment that “string theorists” are no longer doing string theory per se, though they keep identifying themselves as string theorists. As Peter Woit puts it, citing the 2022 “Strings” conference: “one thing that stands out is that the string theory community has almost completely stopped doing string theory.”; and, “[presentations’ titles] make very clear what the string theory community has found to replace string theory: black holes” (Woit, 2022, <https://www.math.columbia.edu/~woit/wordpress/?p=12981>);



(a) Transfers across all research areas.

(b) Flows of attention directed towards “dark matter”.

Figure 4: (a) **Aggregate transfers of attention across research areas, between 2000-2009 (to the left) and 2015-2019 (to the right)**. Widths of flows are proportional to $\sum_a X_{ak}\theta_{akk'}$. Insignificant transfers (that happen less than expected by chance alone assuming uniform mixing) are transparent. (b) For purposes of clarity, the same figure is repeated to the right, highlighting only the flows directed towards “Dark matter”.

3.1 The structuring effect of learning costs on scientists’ behavior

As research priorities change, research efforts must be efficiently reallocated among scientists, given their prior knowledge. Here, we leverage the similarity between our model of scientists’ trajectories and **Optimal Transport (OT)** [56, 57] to formulate the observed behavior in economic terms and demonstrate that the shifts of attention are structured by the minimization of “learning costs”, thus providing a first-order explanation of the aggregate patterns of change in high energy physicists’ research interests. Optimal Transport is a mathematical framework first introduced by Gaspard Monge for finding optimal ways of displacing piles of sand in a military context [12], and later refined by Leonid Kantorovitch in the context of economic planning [13].

Let $\mathbf{x} = \sum_a \mathbf{x}_a$ be the cohort’s initial distribution of attention across research areas (summing over each author a), and $\mathbf{y} = \sum_a \mathbf{y}_a$ the late distribution. Let us further assume that \mathbf{x} and \mathbf{y} can be considered “fixed” by the institutions (scientific leadership, laboratories, funding agencies, etc.) that define scientific priorities throughout time. In order to achieve the distribution of research efforts \mathbf{y} given the previous distribution \mathbf{x} , some scientists must redirect their attention away from certain research areas (for which $y_k < x_k$) and towards more pressing ones (for which $y_k > x_k$). What is the most efficient way to reallocate research efforts and achieve the transition from \mathbf{x} to \mathbf{y} ? Intuitively, research areas should be assigned to scientists in a way that requires as few of them as possible to acquire new knowledge – in other words, in a way that minimizes learning costs, given the way knowledge is distributed among individuals. This, we show, can be framed as an Optimal Transport problem. Let $T_{kk'}$ be the coupling matrix that encodes how much attention has been shifted from one research area k to a research area k' . There are two constraints on $T_{kk'}$: $\sum_{k'} T_{kk'} = x_k$, and $\sum_k T_{kk'} = y_{k'}$. These constraints encode the need to adapt (since $\mathbf{y} \neq \mathbf{x}$). But T must also minimize learning costs (C), which we assume to be linear in T for simplicity¹⁶: $C = \sum_{k,k'} T_{kk'} C_{kk'}$, where $C_{kk'}$ is

¹⁶Roughly speaking, this linear assumption entails that scientists from a given research area are equally able to shift attention to another research area. In practice, some scientists have more abilities to switch to a given research area, and the cost will increase non-linearly as more and more scientists are required to make the transition, including those less able to make the

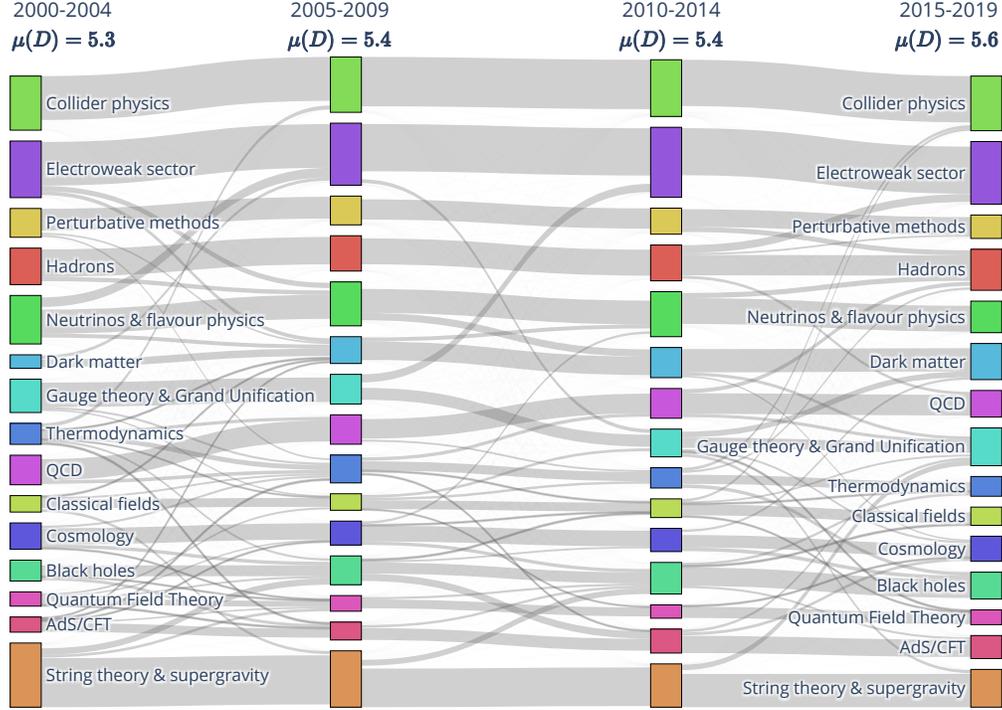


Figure 5: **Aggregate transfers of attention throughout the years 2000 to 2019, considering four time-intervals of five years each.** $\mu(D)$ is the average portfolio diversity of the cohort during each time-period.

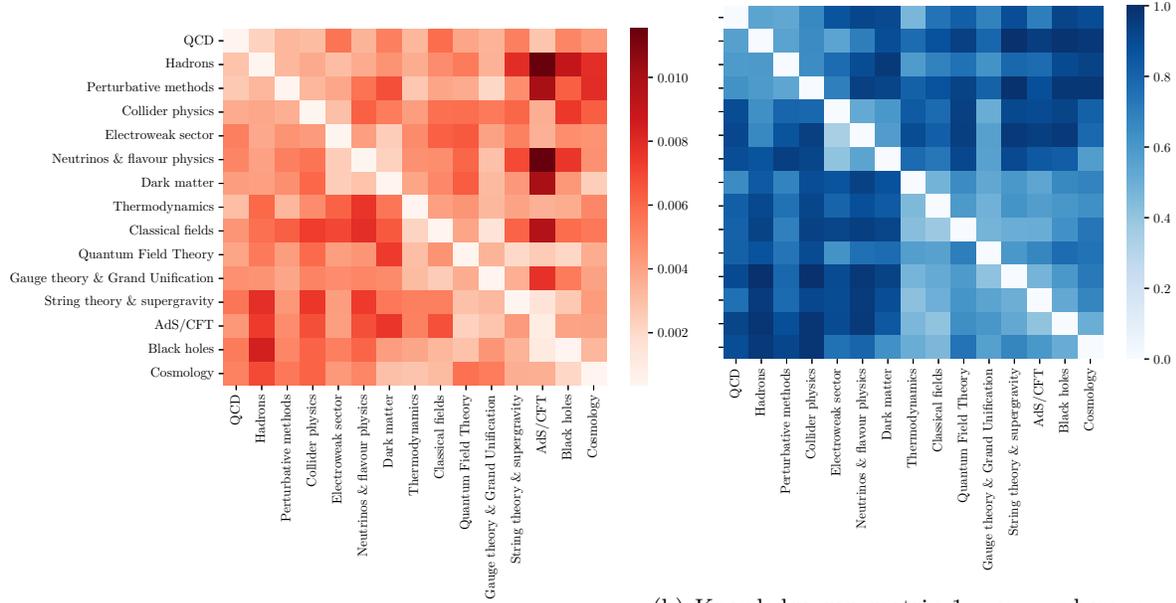
a cost matrix. The problem of finding the couplings $T_{kk'}$ that minimize the “transportation” costs (given a cost matrix $C_{kk'}$ and the constraints on $T_{kk'}$) is an instance of Optimal Transport problem [14]. Typical instances of **OT** problems include how to efficiently transport (say, ore from mines to factories) or the optimal assignment of workers to firms [60]. In our case, the couplings are known (they were recovered by the model from Section 2.4), and we want to infer the underlying cost matrix that these couplings minimize. The transfers from a research area k to k' for each individual are simply $x_{ak}\theta_{akk'}$. Summing over individuals yields the coupling matrix $T_{kk'} = \sum_a x_{ak}\theta_{akk'}$, which measures how much attention was shifted away from k and toward k' at the cohort level.

The problem of recovering the cost matrix $C_{kk'}$ given the couplings is an Inverse Optimal Transport problem. Below, this problem is solved using the probabilistic method proposed in [17]¹⁷. This method requires to put a prior on $P(C_{kk'})$ – indeed, infinitely many cost matrices yield the same optimum, and priors are needed to lift this degeneracy. Following [17], we consider a prior such that $\sum_{k,k'} C_{kk'} = \text{cst}$. We assume that $\mathbb{E}(C_{kk'}) \propto \text{softmax}(\beta \times (1 - \nu_{kk'}))$ ¹⁸, where $\nu_{kk'}$ is the fraction of physicists who already held expertise in k' among those who already held expertise in k , and $\beta \sim \mathcal{N}(0, 1)$ is the effect of learning costs on C_{kk} . If $\nu_{kk'} \sim 1$ (i.e. $1 - \nu_{kk'} \sim 0$), then shifting attention from k to k' does not entail the acquisition of additional knowledge, and $C_{kk'}$ should be almost zero. If $\nu_{kk'} = 0$ (i.e. $1 - \nu_{kk'} = 1$), any scientist shifting from k toward k' must acquire new knowledge, and the cost should be maximal. If actual behaviors *do* involve the minimization of learning costs, we should observe a strong correlation between $C_{kk'}$ and the “knowledge gap” ($1 - \nu_{kk'}$).

switch.

¹⁷It should be noted that the approach by Chiu et al. [17] does not entail the assumption that scientists’ behavior is perfectly minimizing the cost matrix. Indeed, the optimization problem they consider includes an entropic regularization term; while this term is often introduced for numerical reasons, in the case of human behavior, it can be taken to represent inefficiencies and random deviations from the optimum behavior [61].

¹⁸This prior has the merit of simplicity – it is a simple generalized linear model with the desired support.



(a) Cost matrix $C_{kk'}$. The cost matrix represents the cost of shifting a unit of cohort’s research attention from k to k' .

(b) Knowledge gap matrix $1 - \nu_{kk'}$, where $\nu_{kk'}$ is the fraction of physicists with expertise in the horizontal research area (k) who also hold expertise in the vertical research area (k').

Figure 6: **Learning costs and knowledge gap.** The cost of shifting attention from one area to another (left plot) is in great part determined by how frequently individuals hold knowledge in both areas (right plot).

Using the “MetroMC” algorithm proposed in [17], we empirically recover $C_{kk'}$, the underlying cost matrix, as shown in Figure 6a. We find a strong correlation with the “knowledge gap” (shown in Figure 6): $R(C_{kk'}, 1 - \nu_{kk'}) = 0.76$, such that $R^2 = 0.58$ (Figure 17, see also Appendix A.6). Using a finer-grained temporal segmentation (2000-2004→2005-2009), the resulting correlation is similar ($R^2 = 0.62$). The empirical cost of shifting research efforts from one research area to another is therefore shaped by learning costs. The derivation of $C_{kk'}$ is potentially useful: one could in principle predict aggregate transfers of attention given a counterfactual target distribution of research efforts \mathbf{y} using optimal transport and plugging-in $C_{kk'}$ as the cost matrix [57].

Under changing circumstances, research efforts must be reallocated efficiently. Scientific norms and institutions must address an Optimal Transport problem by providing incentives for scientists to conform to new research imperatives, in a way that factors “learning costs”. Consequently, shifts between research areas which entail the acquisition of new knowledge (layering) must be less likely than those which can take advantage of prior knowledge (conversion). In the case of HEP, it does seem that adaptive patterns are structured by this OT problem. Interestingly, the matrices in Figure 6a and 6b feature blocks (indicative of an underlying hierarchical structure), such that it is easier for scientists to migrate within than across these blocks. While these observation characterizes the cohort’s behavior, drivers of differences among individuals are considered next.

3.2 Individual behavior and the effect of capital

3.2.1 Magnitude of change and capital

We propose a change score c_a that measures how much the research agenda of a scientist has changed between the two time periods under consideration, defined as the total variation distance between their initial and late research portfolios:

$$c_a \equiv d_{\text{TV}}(\mathbf{x}_a, \mathbf{y}_a) = \frac{1}{2} \sum_k |y_{ak} - x_{ak}| \quad (12)$$

This measure of change is naturally motivated by Optimal Transport: it is the minimal cost of transporting \mathbf{x}_a to \mathbf{y}_a if the cost matrix has zeros on the diagonal and ones everywhere else ($C_{kk} = 0$ and $C_{kk'} = 1$ for $k \neq k'$). This measure, however, weighs migrations across different research areas identically, regardless of their cognitive proximity. We will therefore also consider a second measure of change, the “cognitive distance” (d_a), defined as the minimum cost of transporting \mathbf{x}_a to \mathbf{y}_a [14] induced by the cost matrix empirically recovered (see Figure 6a) using the Inverse Optimal Transport approach described in the previous section. Another interesting aspect of Optimal Transport is indeed its ability to provide “distances” between distributions that emphasize certain costs in particular ¹⁹.

c_a is comprised between 0 (if research attention has remained identically distributed) and 1 (if the research agenda has been entirely redistributed). Large values of c_a are rare, with 50% of authors lying between 0.21 and 0.40 (Appendix A.7.1, Figure 18). Although the absolute value of c_a (and d_a) has limited interpretability (it depends on the choice of “cognitive scope” and the duration of the time periods considered, and c_a is never expected to be exactly zero due to random fluctuations and measurement noise), it allows for comparisons between physicists. We evaluate the effect of several factors on c_a : i) the diversity of intellectual capital $D(\mathbf{I}_a)$; ii) the excess diversity of social capital $D^*(\mathbf{S}_a)$; and iii) the magnitude of social capital (“power”). We also consider the effect of iv) affiliation stability, represented by a binary variable s_a ($s_a = 1$ if scientist a has at least one affiliation that spans the whole time range considered, and $s_a = 0$ otherwise), the effect of v) academic age, and productivity, estimated from vi) all papers and vii) solo-authored papers. We perform a linear regression of c_a as a function of these variables, adjusting for $Z_a = \arg \max_k x_{ak}$, i.e. physicists’ primary research area over the years 2000 to 2009 (see the model specification in Appendix A.7.1).

The results are shown in Figure 7. The diversity of intellectual capital has a significant positive effect: physicists with resources in many areas tend to revise their research agenda more. There is also evidence of a small but positive effect of diversity of social capital on the magnitude of changes in scientists’ research focus (interpreting these results in terms of Optimal Transport, we might say that social capital helps overcome cognitive learning costs). However, the *magnitude* of social capital, “power”, has a negative direct effect on change. In other words, “power” is associated with stability, and “diversity” with change. It is noteworthy that these dimensions of social capital have opposite effects. More senior physicists are more conservative toward their research agenda, possibly because they experience less incentive to “adapt”. This comes in contrast with [6]. The difference could stem in the specificity of high-energy physics, and in the fact that change is measured in linguistics terms (instead of relying on citation patterns) at a rather coarse-grained scale in the present paper. There is no discernible effect of affiliation stability after adjusting for academic age. However, affiliation data is a bit noisy, and this could have the consequence of underestimating the effect of institutional stability relative to that of academic age. Finally, productivity (in co-authored papers) is associated with stability. Overall, entrenchment (age, power, productivity, specialization) all drive stability and conservatism.

Unsurprisingly, both research areas that have shrunk considerably have a significant positive effect on migration scores. “Collider physics” and “dark matter”, on the other hand, have a negative effect on the magnitude of change. All effects combined, physicists whose primary category is “Collider physics” are the most conservative, with an average change score 24% lower than the rest of the cohort; the long time-scales of collider experiments provide stable opportunities to physicists in that area [62, p. 138], thus promoting conservatism. The variance explained remains low ($R^2 = 0.17$): these factors only partially explain differences between individuals.

c_a neglects the cognitive gap between research areas. Using our alternative measure of change that takes into account cognitive distance (d_a), the diversity of intellectual and social capital has slightly larger effects.

¹⁹Given a cost matrix $C_{kk'}$, we can define a measure of the gap from one distribution \mathbf{x} to another distribution \mathbf{y} , as the minimum cost of displacing \mathbf{x} to \mathbf{y} :

$$d(\mathbf{x}, \mathbf{y}) = \min_{\theta_{kk'}, \theta \mathbf{1} = \mathbf{1}, \theta^T \mathbf{x} = \mathbf{y}} \sum_{kk'} x_k \theta_{kk'} C_{kk'} \quad (13)$$

If the cost matrix meets certain properties (such as symmetry), then d is a distance. See [14] for more on the metric properties of OT.

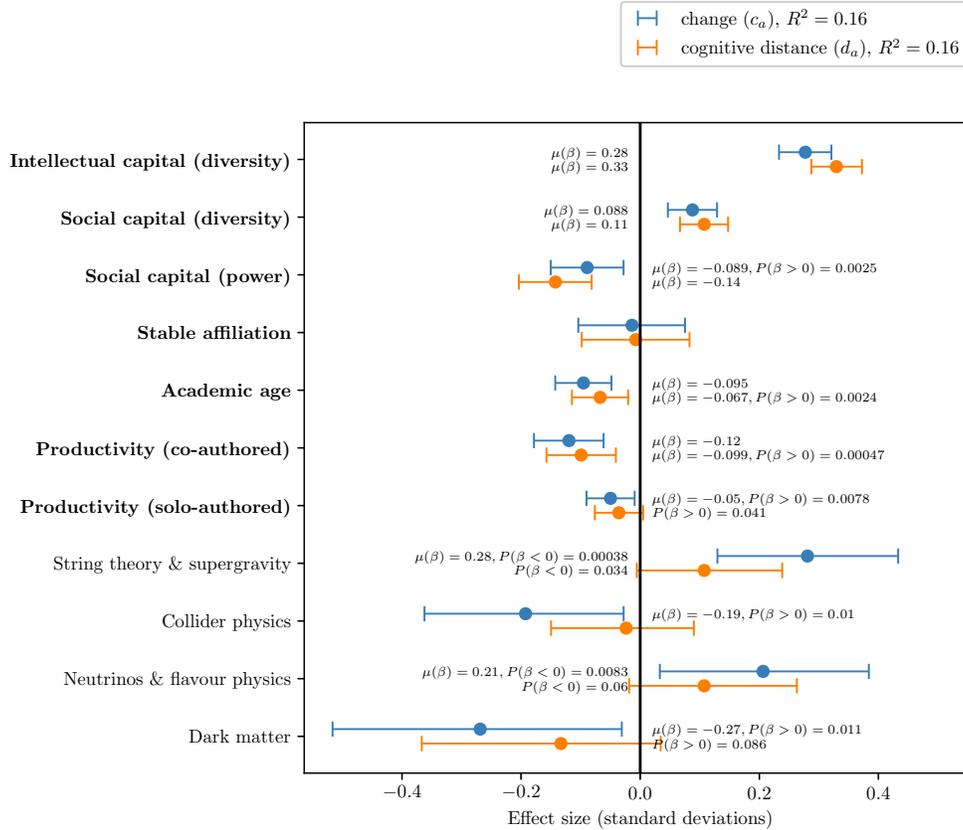


Figure 7: **Effect of intellectual capital, social capital, and institutional stability the magnitude of changes in high-energy physicists’ interests.** Change (c_a) is the total variation distance and weighs migrations across research areas equally. The cognitive distance (d_a) gives more weight to migrations across cognitively distant research area. Diversity is evaluated as the exponential of the Shannon entropy of the distribution of intellectual and social capital across research areas. The effect of the primary research area during the initial period (2000-2009) was also included as a control. Three research areas have a significant (95% CL) effect. Error bars indicate 95% credible intervals; $\mu(\beta)$ denotes the mean posterior effect size. $R^2 = 0.17$. Continuous variables (diversity, power and c_a) are standardized.

The robustness of these results is assessed by using different operationalizations of diversity and power in Appendix A.7.3), Table 3. Most findings are stable, except that “brokerage” (unlike the magnitude of social capital, i.e. degree centrality) has no *direct* effect on change (beyond the effect of productivity resulting from co-authored publications).

In addition, in order to exclude the influence of direct collaborations on change, we conducted a second comparative analysis including only scientists’ first-authored and last-authored publications in their research portfolios²⁰. The effect of the diversity of intellectual and social capital is stable; however, the direct effect of power is reduced (Appendix A.7.3, Tables 5 and 6). Moreover, the analysis was reproduced across different choices of temporal segmentation and using different topic models. The previous findings remain stable.

3.2.2 Diversification versus concentration

Research portfolios can be altered by two opposite strategies. One is “diversification”, i.e. the addition of new research areas. Another is “concentration”, i.e. the desertion of research areas to focus increasingly on others (Figure 1 illustrates how this can happen, whether via “conversion” or “displacement”). Figure 8

²⁰It is important to emphasize that alphabetical ordering is still very prevalent in this field [63], and therefore this strategy does not address the issue entirely.

shows the effect of the same factors as above on the probability that physicists have i) entered at least one new research area in between the two periods or ii) exited one research area (model description provided in Appendix A.7.2)²¹.

The diversity of intellectual capital has a strong positive effect on the probability of exiting a research area; intuitively, scientists with diverse interests can afford to disengage from certain research areas even if this implies to abandon maladaptive prior knowledge (“displacement”). Excess diversity of social capital increases the probability of entering new research areas, but has no discernible effect on the probability of exiting research areas. In contrast, there is some evidence that “power” decreases the probability of leaving a research area. Figure 8 shows the direct effect of power (controlling for productivity due to co-authored papers), which does not pass the 95% significance test; however, the *total* effect of power²² on the probability of leaving a research area is significantly negative ($\mu(\beta) = -0.24, P(\beta > 0) < 10^{-4}$). Presumably, having many collaborators allows scientists to remain committed to many research areas with minimal personal investment, thus stabilizing their research agendas. Moreover, “power” also decreases the probability of entering new research areas; this suggests that powerful scientists have less incentives to invest resources in new topics. String theory & supergravity has a significant positive effect of both entering and exiting research areas (suggesting physicists frequently “replaced” it with another topic). “Dark matter” has a negative effect on the probability of entering a new research area, possibly because scientists with prior commitment to this research area had less incentive to diversify their research portfolio and more incentive to focus increasingly on this topic given its success over the following years. These conclusions hold as alternative measures of diversity and power are considered (A.7.3, Table 4).

Again, we ran a second analysis considering only scientists’ first-authored and last-authored publications. The effects of intellectual and social diversity remain stable. However, the effect of productivity resulting from co-authored papers no longer has an effect on the probability of exiting a research area when only first- and last-authored papers are included in physicists’ research portfolios (8, Appendix A.7.3.). The results are generally stable across different topic models and temporal segmentations, except for i) the effect of academic age on the probability of exiting a research area, which is not consistently 95% CL significant; and ii) the effect of the diversity of social capital on the probability of entering a new research area, which is zero in one particular configuration.

3.2.3 Why diversity promotes change

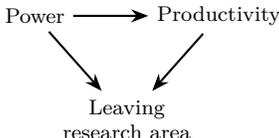
Access to diverse cognitive resources is associated with change. To see why, it is insightful to look into how the concentration of scientists’ intellectual capital in each research area k ($\mathbf{I}_a = (I_{ak})$) affects their trajectories. In the model introduced in Section 2.4, the diagonal coefficients of the γ matrix measure the effect of having intellectual resources in a certain area on the commitment to this area. As shown in Figure 9a, most coefficients on the diagonal are significantly positive: physicists with a strong specialization in a research area tend to preserve their specialization in this area.

The effect of social capital on transfers across research areas ($\delta_{kk'}$, Eq. 6), is shown in Figure 9b. Statistically significant effects are always positive: scientists tend to redirect attention to research areas in which they have more collaborators involved, in line with a very recent observation by Venturini et al. [64].

Moreover, we find a strong correlation between the effect of social capital ($\delta_{kk'}$), and $\nu_{kk'} \in [0, 1]$, the fraction of physicists with more expertise than average in a research area k' among those who have more expertise than average in k (see Eq. (9), Section 2.4). $\delta_{kk'}$ decreases by 1.5 unit of standard deviation on average as $\nu_{kk'}$ goes from 0 (nobody holds expertise in both research areas k and k') to 1 (everyone with expertise in k has expertise in k'). Social capital plays a more important role in shifts between cognitively distant research areas, in line with a previous finding by Tripodi et al. [7]. These general patterns (the

²¹A research area k is considered “entered” by a when $x_{ak} < \frac{1}{N} \sum_{a'} x_{a'k}$ and $y_{ak} > \frac{1}{N} \sum_{a'} y_{a'k}$; conversely, a research area is considered exited when $x_{ak} > \frac{1}{N} \sum_{a'} x_{a'k}$ and $y_{ak} < \frac{1}{N} \sum_{a'} y_{a'k}$.

²²That is, assuming the following causal structure:



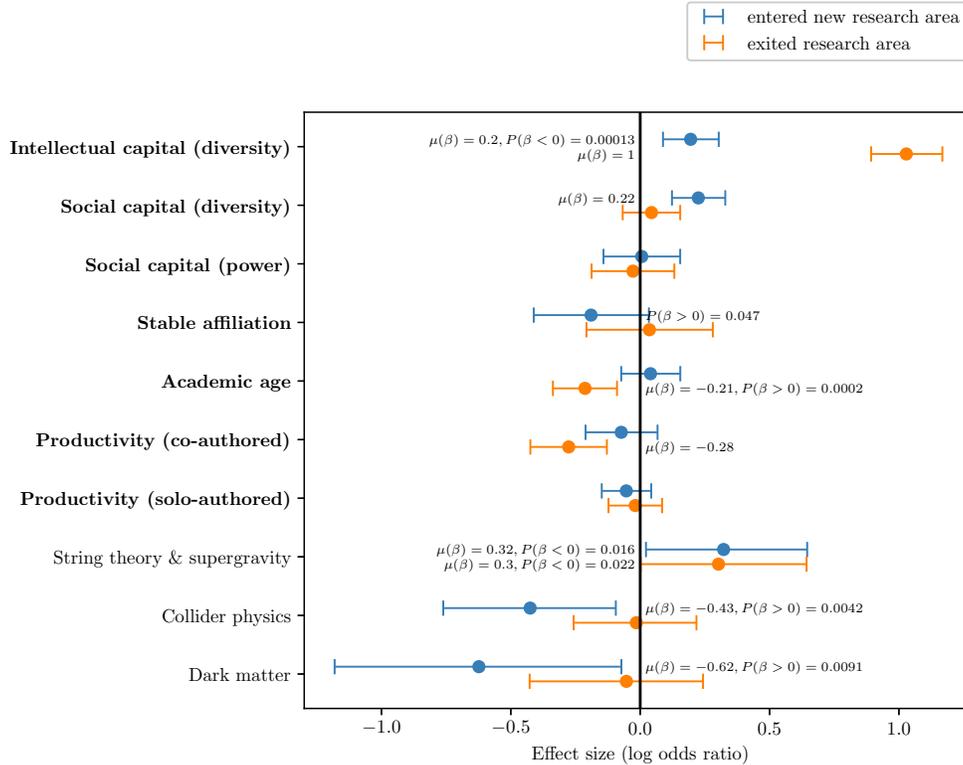


Figure 8: **Effect of capital and institutional stability on the probability of entering new research areas or exiting previously explored ones.** Bars indicate 95% credible intervals. $\mu(\beta)$ denotes the mean posterior effect size. Continuous variables (diversity and power) are standardized.

association between the concentration of intellectual capital in one area and the commitment to this area, and the increasing effect of social capital with cognitive dissimilarity) are insensitive to the temporal segmentation in place (see Appendix A.9 for similar Figures based on a different segmentation).

4 Discussion

This paper addressed the conflict between specialization and adaptability in science. To this end, an unsupervised Bayesian approach was developed, based on the idea that transformations in the scientific landscape prompt scientists to efficiently repurpose their prior knowledge. The model simultaneously measures transfers of attention across research areas and the effect of various variables on the evolution of scientists' research portfolios, in particular intellectual and social capital.

The model was applied to a cohort of $N = 2\,108$ high-energy physicists between the years 2000 and 2019. At the macroscopic level, it reveals the decline of neutrinos physics due to migrations towards the electroweak sector (explored at the LHC) and more importantly towards dark matter. Similarly, many physicists have shifted resources from the electroweak sector towards dark matter. Moreover, string theory & supergravity has started to disintegrate into black holes and AdS/CFT research. The cohort has therefore responded to new experimental opportunities as well as to theoretical developments in quantum gravity.

Then, leveraging the connection between our model and **Optimal Transport (OT)**, we showed that the reallocation of research efforts among scientists is shaped by learning costs. Indeed, under changing circumstances, scientific institutions must address an **OT** problem by efficiently reallocating research efforts in a way that balances learning costs and the imperative to adapt to new circumstances. This has the effect of enhancing the utility of the scientific capital disseminated amongst scientists as the perceived payoff of certain research areas change. **OT** structurally explains path dependency, as individuals experiencing pressures to

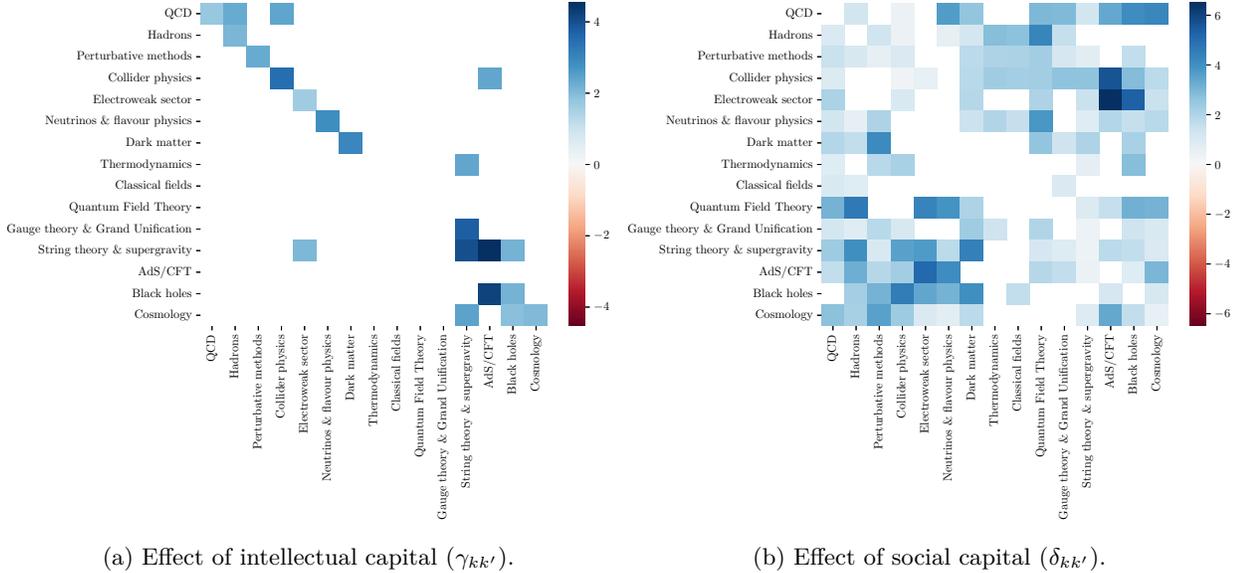


Figure 9: **Effect of intellectual and social capital on transfers across research areas.** Rows represent research areas of origin and columns represent target areas. Effect sizes are expressed in log-odds ($\log [\theta_{akk'}/(1 - \theta_{akk'})]$, where $\theta_{akk'}$ is the fraction of attention to k redirected to k') per unit of intellectual or social capital in the target research area k' . Effects that are not significant (at the 95% credible level) are displayed in white for purposes of clarity.

adapt seize the nearest opportunity available to them. **OT** is also methodologically useful: Inverse **OT** allows one to derive cost functions from observed behavior, and thus offers a potential way to better connect empirical data with evolutionary agent-based models of scientists’ behavior that postulate latent utility functions [16]. Moreover, it has the potential of informing policy by identifying potential bottlenecks if research efforts were to be reallocated in certain ways (as one can use **OT** to estimate the “cost” of various counterfactual scenarios). Overall, the **OT** approach illustrates that the adaptability of epistemic communities is constrained by how knowledge is distributed among individuals (specialization).

The longitudinal comparative analysis of physicists’ trajectories revealed that the diversity of intellectual and social capital is positively associated with change: diversity promotes adaptability under new circumstances, and therefore diversifying research portfolios is a reasonable strategy when the future is uncertain. However, enforcing diversity can lead to suboptimal allocations of research efforts under stable circumstances [30]. Differences among research areas are found: physicists expert in particle colliders have remained particularly conservative, possibly because they have secured long-term research opportunities (thanks to very large investments in particle accelerators like the **LHC**). There is also evidence that physicists specialized in dark matter have been consolidating their specialization, presumably because the increasing popularity of the topic encouraged them to double down their investment in this research area. Higher concentrations of intellectual capital in certain research areas generate stronger commitment towards these areas; therefore, specialized scientists are more at risk of being trapped in a sunk cost fallacy as their expertise becomes unsuitable for new circumstances. However, specialized scientists can offset the risks associated with specialization by diversifying their social network. This raises the possibility of free-riding, as scientists are encouraged to focus on what seems most promising at the time and let their peers take the risk of exploring alternatives until their value is established [65]. Additionally, social capital plays an increasingly important role as scientists expand their research agenda further beyond their specialization, as observed by [7], suggesting that collaborations are crucial in overcoming cognitive barriers between research areas. Unlike diversity, “power” is associated with more stable research interests: presumably, cooperation can safeguard individuals from adaptive pressures, and most importantly minimizes the cost of remaining invested in certain areas.

We have described renewal strategies of research portfolios according to a typology of incremental change developed in the context of historical institutionalism, which has been shown to account for how organizations

like DESY and SLAC have transitioned from **HEP** to multi-purpose photon science. Adaptation strategies include the “conversion” of prior knowledge to new purposes, the “layering” of additional research interests via the acquisition of new knowledge, or the “displacement” of former research commitments resulting in a loss of knowledge. The connection with institutional change stems from the fact that institutions can face a challenge similar to that experienced by specialized scientists confronted with transformations in their scientific landscape, in that they too must sometimes adapt and redirect accumulated capital in directions that may not have been foreseen at the time of their “design”. The **LHC** itself has evolved through similar processes of gradual repurposing of prior infrastructure, including the accelerator’s tunnels [66]: adaptation prompts individuals (and collectives) to efficiently repurpose capital previously accumulated in different forms (cultural, social, economic, etc.). Following [11], we conclude that a better understanding of collective adaptation benefits from the pooling of diverse insights: while studies of institutional change have documented strategies of gradual adaptation that progressively leverage and repurpose accumulated capital when change is difficult, the literature on cultural evolvability stresses the critical roles of diversity and social learning. In return, this work provides an empirical contribution to the literature that treats science as a cultural evolutionary system [16].

Throughout the paper, various methodological limitations were raised. First, the requirement on the quantity of authors’ publications makes the cohort atypical. A second issue, already noted by Gieryn [18], lies in the arbitrariness in the choice of temporal and cognitive scopes for measuring change. In this paper, the choice in temporal segmentation was driven by the time-scale associated with the transformations in the landscape of experimental opportunities. For shorter time-scales, the very notion of research portfolio – as operationalized in this paper – may break down and lose any predictive force. As per the cognitive scope, previous works (e.g. [4, 5, 7]) have typically relied on the hierarchical **PACS** classification of physics literature, such that cognitive scales were imposed. For this work, the literature was divided into 15 topics that captured the features of change in the epistemic landscape discussed above, i.e. the rise of new probes of the cosmos, but other scales could have been considered. Moreover, although the unsupervised topic model approach is arguably a better proxy of cognitive change, it introduces noise which could in part explain the low predictive power of the models of change in scientists’ research agendas. Additionally, the topic model was trained on the entire time range covered by the analysis. Changes in the relationships between topics and in their own vocabulary distributions are not considered, even though they constitute another interesting linguistic dimension of adaptive patterns that would deserve further investigation. Moreover, the cost of shifting from one research area to another, is itself time-varying quantity in reality. Finally, one must be cautious before drawing strong causal conclusions from these findings. While the causal pathway “power → collaborations → stabilization of research interests” seems to be a reasonable interpretation of the results, the relationship between diversity and change is less clear; they could both be confounded by a latent trait specific to certain researchers (e.g. the “explorers” in [67]). Moreover, the sample size ($N = 2108$) is insufficient to explore sophisticated interactions or potential moderators in the comparative analysis.

Acronyms

APS American Physical Society.

HEP High-Energy Physics.

LHC Large Hadron Collider.

OT Optimal Transport.

PACS Physics and Astronomy Classification Scheme®.

Declarations

Availability of data and material The data and code used for this publication is available at <https://github.com/lucasgautheron/specialization-adaptation/>. The repository will be archived on Zenodo upon publication.

Competing Interests The author declares no competing interest.

Funding This work has received support from the G-Research Grant for Early-Career Researchers.

Authors' contributions The sole author of the paper completed all of the work.

Acknowledgements I thank Thomas Heinze, Radin Dardashti, and Elisa Omodei for useful discussions; Marc Santolini for rich feedback; Georges Ricci for raising my interest in Optimal Transport; Marco Schirone for providing useful references; the LATTICE lab in Montrouge for the opportunity to discuss this work; and Elizabeth Zanghi for corrections.

References

- [1] T. S. Kuhn. “The Essential Tension: Tradition and Innovation in Scientific Research”. In: *The Essential Tension*. University of Chicago Press, 1977, pp. 225–239. DOI: [doi:10.7208/9780226217239-010](https://doi.org/10.7208/9780226217239-010).
- [2] J. G. March. “Exploration and Exploitation in Organizational Learning”. In: *Organization Science* 2.1 (Feb. 1991), pp. 71–87. DOI: [10.1287/orsc.2.1.71](https://doi.org/10.1287/orsc.2.1.71).
- [3] J. G. Foster, A. Rzhetsky, and J. A. Evans. “Tradition and Innovation in Scientists’ Research Strategies”. In: *American Sociological Review* 80.5 (Sept. 2015), pp. 875–908. DOI: [10.1177/0003122415601618](https://doi.org/10.1177/0003122415601618).
- [4] T. Jia, D. Wang, and B. K. Szymanski. “Quantifying patterns of research-interest evolution”. In: *Nature Human Behaviour* 1.4 (Mar. 2017). DOI: [10.1038/s41562-017-0078](https://doi.org/10.1038/s41562-017-0078).
- [5] A. Aleta, S. Meloni, N. Perra, and Y. Moreno. “Explore with caution: mapping the evolution of scientific interest in physics”. In: *EPJ Data Science* 8.1 (Sept. 2019). DOI: [10.1140/epjds/s13688-019-0205-9](https://doi.org/10.1140/epjds/s13688-019-0205-9).
- [6] A. Zeng, Z. Shen, J. Zhou, Y. Fan, Z. Di, Y. Wang, H. E. Stanley, and S. Havlin. “Increasing trend of scientists to switch between topics”. In: *Nature Communications* 10.1 (July 2019). DOI: [10.1038/s41467-019-11401-8](https://doi.org/10.1038/s41467-019-11401-8).
- [7] G. Tripodi, F. Chiaromonte, and F. Lillo. “Knowledge and social relatedness shape research portfolio diversification”. In: *Scientific Reports* 10.1 (Aug. 2020). DOI: [10.1038/s41598-020-71009-7](https://doi.org/10.1038/s41598-020-71009-7).
- [8] C. K. Singh, L. Tupikina, F. Lécuyer, M. Starnini, and M. Santolini. “Charting mobility patterns in the scientific knowledge landscape”. In: *EPJ Data Science* 13.1 (Feb. 2024). DOI: [10.1140/epjds/s13688-024-00451-8](https://doi.org/10.1140/epjds/s13688-024-00451-8).
- [9] F. Liu, S. Zhang, and H. Xia. “Science as exploration in a knowledge landscape: tracing hotspots or seeking opportunity?” In: *EPJ Data Science* 13.1 (2024), p. 27.
- [10] J. H. Fujimura. “The Molecular Biological Bandwagon in Cancer Research: Where Social Worlds Meet”. In: *Social Problems* 35.3 (June 1988), pp. 261–283. DOI: [10.2307/800622](https://doi.org/10.2307/800622).
- [11] M. Galesic et al. “Beyond collective intelligence: Collective adaptation”. In: *Journal of The Royal Society Interface* 20.200 (Mar. 2023). DOI: [10.1098/rsif.2022.0736](https://doi.org/10.1098/rsif.2022.0736).
- [12] G. Monge. “Mémoire sur la théorie des déblais et des remblais”. In: *Mem. Math. Phys. Acad. Royale Sci.* (1781), pp. 666–704.
- [13] L. Kantorovich. “On the Translocation of Masses”. In: *C.R. (Doklady) Acad. Sci. URSS (N.S.)* 37 (1942), pp. 199–201.
- [14] G. Peyré and M. Cuturi. “Computational Optimal Transport: With Applications to Data Science”. In: *Foundations and Trends in Machine Learning* 11.56 (2019), pp. 355–607. DOI: [10.1561/22000000073](https://doi.org/10.1561/22000000073).
- [15] A. Abbasi, R. T. Wigand, and L. Hossain. “Measuring social capital through network analysis and its influence on individual performance”. In: *Library & Information Science Research* 36.1 (Jan. 2014), pp. 66–73. DOI: [10.1016/j.lisr.2013.08.001](https://doi.org/10.1016/j.lisr.2013.08.001).
- [16] J. Wu, C. O’Connor, and P. E. Smaldino. “The Cultural Evolution of Science”. In: *The Oxford Handbook of Cultural Evolution*. Oxford University Press, Sept. 2023. DOI: [10.1093/oxfordhb/9780198869252.013.78](https://doi.org/10.1093/oxfordhb/9780198869252.013.78).

- [17] W.-T. Chiu, P. Wang, and P. Shafto. “Discrete Probabilistic Inverse Optimal Transport”. In: *Proceedings of the 39th International Conference on Machine Learning*. Ed. by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato. Vol. 162. Proceedings of Machine Learning Research. PMLR, July 2022, pp. 3925–3946.
- [18] T. F. Gieryn. “Problem Retention and Problem Change in Science”. In: *Sociological Inquiry* 48.3-4 (July 1978), pp. 96–115. DOI: [10.1111/j.1475-682x.1978.tb00820.x](https://doi.org/10.1111/j.1475-682x.1978.tb00820.x).
- [19] F. Battiston, F. Musciotto, D. Wang, A.-L. Barabási, M. Szell, and R. Sinatra. “Taking census of physics”. In: *Nature Reviews Physics* 1.1 (Jan. 2019), pp. 89–97. DOI: [10.1038/s42254-018-0005-3](https://doi.org/10.1038/s42254-018-0005-3).
- [20] P. Bourdieu. “Le capital social”. In: *Actes de la Recherche en Sciences Sociales* 31.1 (1980), pp. 2–3.
- [21] P. Bourdieu. “The Forms of Capital”. In: *Handbook of Theory and Research for the Sociology of Education*. Ed. by J. Richardson. Greenwood, 1986, pp. 241–258.
- [22] M. Mulkay. “Conceptual Displacement and Migration in Science: A Prefatory Paper”. In: *Science Studies* 4.3 (July 1974), pp. 205–234. DOI: [10.1177/030631277400400301](https://doi.org/10.1177/030631277400400301).
- [23] D. A. Schon. *Displacement of Concepts*. 1st. Routledge, 1963. DOI: [10.4324/9781315014111](https://doi.org/10.4324/9781315014111).
- [24] J. Mahoney and K. Thelen. “A Theory of Gradual Institutional Change”. In: *Explaining Institutional Change: Ambiguity, Agency, and Power*. Cambridge University Press, 2009, pp. 1–37. DOI: [10.1017/CB09780511806414.003](https://doi.org/10.1017/CB09780511806414.003).
- [25] O. Hallonsten and T. Heinze. “From particle physics to photon science: Multi-dimensional and multi-level renewal at DESY and SLAC”. In: *Science and Public Policy* 40.5 (Mar. 2013), pp. 591–603. DOI: [10.1093/scipol/sct009](https://doi.org/10.1093/scipol/sct009).
- [26] O. Hallonsten and T. Heinze. “Formation and expansion of a new organizational field in experimental science”. In: *Science and Public Policy* 42.6 (2015), pp. 841–854.
- [27] T. Heinze, O. Hallonsten, and S. Heinecke. “Turning the Ship: The Transformation of DESY, 1993–2009”. In: *Physics in Perspective* 19.4 (Sept. 2017), pp. 424–451. DOI: [10.1007/s00016-017-0209-4](https://doi.org/10.1007/s00016-017-0209-4).
- [28] T. Heinze and R. Münch. “Intellektuelle Erneuerung der Forschung durch institutionellen Wandel”. In: *Institutionelle Erneuerungsfähigkeit der Forschung*. VS Verlag für Sozialwissenschaften, 2012, pp. 15–38. DOI: [10.1007/978-3-531-94274-2_2](https://doi.org/10.1007/978-3-531-94274-2_2).
- [29] P. E. Smaldino, C. Moser, A. Pérez Velilla, and M. Werling. “Maintaining Transient Diversity Is a General Principle for Improving Collective Problem Solving”. In: *Perspectives on Psychological Science* 19.2 (June 2023), pp. 454–464. DOI: [10.1177/17456916231180100](https://doi.org/10.1177/17456916231180100).
- [30] R. Schimmelpfennig, L. Razek, E. Schnell, and M. Muthukrishna. “Paradox of diversity in the collective brain”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 377.1843 (Dec. 2021). DOI: [10.1098/rstb.2020.0316](https://doi.org/10.1098/rstb.2020.0316).
- [31] M. Muthukrishna and J. Henrich. “Innovation in the collective brain”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 371.1690 (Mar. 2016), p. 20150192. DOI: [10.1098/rstb.2015.0192](https://doi.org/10.1098/rstb.2015.0192).
- [32] S. Page. *The difference: How the power of diversity creates better groups, firms, schools, and societies-new edition*. Princeton University Press, 2008.
- [33] A. Acerbi, S. Ghirlanda, and M. Enquist. “Old and Young Individuals Role in Cultural Change”. In: *Journal of Artificial Societies and Social Simulation* 15.4 (2012). DOI: [10.18564/jasss.2023](https://doi.org/10.18564/jasss.2023).
- [34] J. Lykken and M. Spiropulu. “Supersymmetry and the Crisis in Physics”. In: *Scientific American* 310.5 (2014), pp. 34–39.
- [35] T. Roser et al. “On the feasibility of future colliders: report of the Snowmass’21 Implementation Task Force”. In: *Journal of Instrumentation* 18.05 (May 2023), P05018. DOI: [10.1088/1748-0221/18/05/p05018](https://doi.org/10.1088/1748-0221/18/05/p05018).
- [36] R. Harlander, J.-P. Martinez, and G. Schieman. “The end of the particle era?” In: *The European Physical Journal H* 48.1 (June 2023). DOI: [10.1140/epjh/s13129-023-00053-4](https://doi.org/10.1140/epjh/s13129-023-00053-4).
- [37] B. Kosyakov. *A farewell to particles*. 2023. DOI: [10.48550/ARXIV.2305.09692](https://doi.org/10.48550/ARXIV.2305.09692).

- [38] M. Shifman. “Musings on the current status of HEP”. In: *Modern Physics Letters A* 35.07 (Feb. 2020), p. 2030003. DOI: [10.1142/s0217732320300037](https://doi.org/10.1142/s0217732320300037).
- [39] B. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, F. Acernese, K. Ackley, and C. Adams. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116.6 (Feb. 2016). DOI: [10.1103/physrevlett.116.061102](https://doi.org/10.1103/physrevlett.116.061102).
- [40] L. Gautheron and E. Omodei. “How research programs come apart: The example of supersymmetry and the disunity of physics”. In: *Quantitative Science Studies* 4.3 (2023), pp. 671–699. DOI: [10.1162/qss_a_00262](https://doi.org/10.1162/qss_a_00262).
- [41] M. Moskovic. *The INSPIRE REST API*. 2021. DOI: [10.5281/ZENODO.5788550](https://doi.org/10.5281/ZENODO.5788550).
- [42] S. Perovi, S. Radovanovi, V. Sikimi, and A. Berber. “Optimal research team composition: data envelopment analysis of Fermilab experiments”. In: *Scientometrics* 108.1 (Apr. 2016), pp. 83–111. DOI: [10.1007/s11192-016-1947-9](https://doi.org/10.1007/s11192-016-1947-9).
- [43] C. Chall, M. King, P. Mättig, and M. Stöltzner. “From a boson to the standard model Higgs: a case study in confirmation and model dynamics”. In: *Synthese* 198.S16 (Apr. 2019), pp. 3779–3811. DOI: [10.1007/s11229-019-02216-7](https://doi.org/10.1007/s11229-019-02216-7).
- [44] A. Strumia. “Gender issues in fundamental physics: A bibliometric analysis”. In: *Quantitative Science Studies* 2.1 (2021), pp. 225–253. DOI: [10.1162/qss_a_00114](https://doi.org/10.1162/qss_a_00114).
- [45] V. Sikimi and S. Radovanovi. “Machine learning in scientific grant review: algorithmically predicting project efficiency in high energy physics”. In: *European Journal for Philosophy of Science* 12.3 (July 2022). DOI: [10.1007/s13194-022-00478-6](https://doi.org/10.1007/s13194-022-00478-6).
- [46] J. P. A. Ioannidis, K. W. Boyack, and R. Klavans. “Estimates of the Continuously Publishing Core in the Scientific Workforce”. In: *PLoS ONE* 9.7 (July 2014). Ed. by L. A. N. Amaral, e101698. DOI: [10.1371/journal.pone.0101698](https://doi.org/10.1371/journal.pone.0101698).
- [47] A. B. Dieng, F. J. R. Ruiz, and D. M. Blei. “Topic Modeling in Embedding Spaces”. In: *Transactions of the Association for Computational Linguistics* 8 (Dec. 2020), pp. 439–453. DOI: [10.1162/tac1_a_00325](https://doi.org/10.1162/tac1_a_00325).
- [48] D. M. Blei, A. Y. Ng, and M. I. Jordan. “Latent Dirichlet Allocation”. In: *J. Mach. Learn. Res.* 3. Jan (Mar. 2003), pp. 993–1022.
- [49] E. Omodei. “Modeling the socio-semantic dynamics of scientific communities”. Thesis. Ecole Normale Supérieure, Dec. 2014.
- [50] T. Mikolov, K. Chen, G. Corrado, and J. Dean. “Efficient Estimation of Word Representations in Vector Space”. In: *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*. Ed. by Y. Bengio and Y. LeCun. 2013.
- [51] M. Schirone. “Field, capital, and habitus: The impact of Pierre Bourdieu on bibliometrics”. In: *Quantitative Science Studies* 4.1 (2023), pp. 186–208. DOI: [10.1162/qss_a_00232](https://doi.org/10.1162/qss_a_00232).
- [52] C. Roth and J.-P. Cointet. “Social and semantic coevolution in knowledge networks”. In: *Social Networks* 32.1 (Jan. 2010), pp. 16–29. DOI: [10.1016/j.socnet.2009.04.005](https://doi.org/10.1016/j.socnet.2009.04.005).
- [53] R. S. Burt. *Brokerage and Closure: An Introduction to Social Capital*. Oxford University Press, 2007. 296 pp.
- [54] M. E. Newman. “Who Is the Best Connected Scientist? A Study of Scientific Coauthorship Networks”. In: *Complex Networks*. Springer, Aug. 2004, pp. 337–370. DOI: [10.1007/978-3-540-44485-5_16](https://doi.org/10.1007/978-3-540-44485-5_16).
- [55] L. Jost. “Entropy and diversity”. In: *Oikos* 113.2 (May 2006), pp. 363–375. DOI: [10.1111/j.2006.0030-1299.14714.x](https://doi.org/10.1111/j.2006.0030-1299.14714.x).
- [56] B. Muzellec, R. Nock, G. Patrini, and F. Nielsen. “Tsallis regularized optimal transport and ecological inference”. In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 31. 1. 2017.
- [57] R. Li, X. Ye, H. Zhou, and H. Zha. “Learning to match via inverse optimal transport”. In: *Journal of machine learning research* 20 (2019).

- [58] N. C. M. Martens and M. King. “Doing More with Less: Dark Matter & Modified Gravity”. In: *Philosophy of Astrophysics*. Springer International Publishing, 2023, pp. 91–107. DOI: [10.1007/978-3-031-26618-8_6](https://doi.org/10.1007/978-3-031-26618-8_6).
- [59] N. C. Martens, M. Á. C. Sahuquillo, E. Scholz, D. Lehmkuhl, and M. Krämer. *Integrating dark matter, modified gravity, and the humanities*. 2022.
- [60] A. Galichon. *Optimal transport methods in economics*. Princeton University Press, 2018.
- [61] A. Dupuy and A. Galichon. “Personality traits and the marriage market”. In: *Journal of Political Economy* 122.6 (2014), pp. 1271–1319.
- [62] P. Galison. *How experiments end*. University of Chicago Press, 1987.
- [63] L. Waltman. “An empirical analysis of the use of alphabetical authorship in scientific publishing”. In: *Journal of Informetrics* 6.4 (2012), pp. 700–711.
- [64] S. Venturini, S. Sikdar, F. Rinaldi, F. Tudisco, and S. Fortunato. “Collaboration and topic switches in science”. In: *Scientific Reports* 14.1 (Jan. 2024). DOI: [10.1038/s41598-024-51606-6](https://doi.org/10.1038/s41598-024-51606-6).
- [65] E. Kummerfeld and K. J. S. Zollman. “Conservatism and the Scientific State of Nature”. In: *The British Journal for the Philosophy of Science* 67.4 (Dec. 2016), pp. 1057–1076. DOI: [10.1093/bjps/axv013](https://doi.org/10.1093/bjps/axv013).
- [66] C. L. Smith. “Genesis of the Large Hadron Collider”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373.2032 (Jan. 2015), p. 20140037. DOI: [10.1098/rsta.2014.0037](https://doi.org/10.1098/rsta.2014.0037).
- [67] C. K. Singh, L. Tupikina, F. Lécuyer, M. Starnini, and M. Santolini. *Charting mobility patterns in the scientific knowledge landscape*. 2023. DOI: [10.48550/ARXIV.2302.13054](https://doi.org/10.48550/ARXIV.2302.13054).
- [68] A. Stirling. “A general framework for analysing diversity in science, technology and society”. In: *Journal of The Royal Society Interface* 4.15 (Feb. 2007), pp. 707–719. DOI: [10.1098/rsif.2007.0213](https://doi.org/10.1098/rsif.2007.0213).
- [69] A. L. Porter, A. S. Cohen, J. D. Roessner, and M. Perreault. “Measuring researcher interdisciplinarity”. In: *Scientometrics* 72.1 (June 2007), pp. 117–147. DOI: [10.1007/s11192-007-1700-5](https://doi.org/10.1007/s11192-007-1700-5).
- [70] E. Leahey, C. M. Beckman, and T. L. Stanko. “Prominent but Less Productive”. In: *Administrative Science Quarterly* 62.1 (Sept. 2016), pp. 105–139. DOI: [10.1177/0001839216665364](https://doi.org/10.1177/0001839216665364).
- [71] S. Rohatgi. *ACL Anthology Corpus with Full Text*. Github. 2022.
- [72] J. Skilling and S. F. Gull. “Bayesian maximum entropy image reconstruction”. In: *Lecture Notes-Monograph Series* (1991), pp. 341–367.
- [73] D. J. C. MacKay and L. C. B. Peto. “A hierarchical Dirichlet language model”. In: *Natural Language Engineering* 1.3 (Sept. 1995), pp. 289–308. DOI: [10.1017/s1351324900000218](https://doi.org/10.1017/s1351324900000218).

A Supplementary material

A.1 Transformations in high-energy physics

Top-25 most cited experiments in HEP, per year

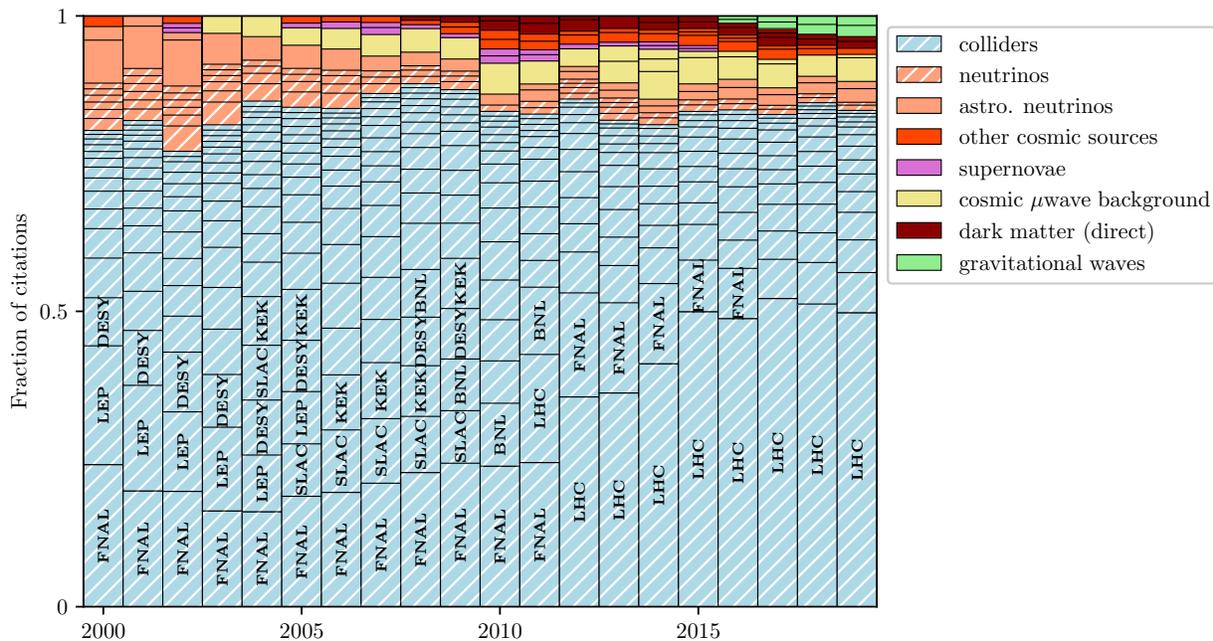


Figure 10: **Share of citations of each of the most cited experiments in high-energy physics literature, between 2000 and 2019.** Hatched rectangles correspond to experiments observing particles produced in colliders or nuclear reactors; other rectangles correspond to observations of phenomena or particles of astrophysical origin. We use citation and experiment data from the Inspire HEP database [41]. The classification of these experiments is our own.

A.2 Cohort characteristics

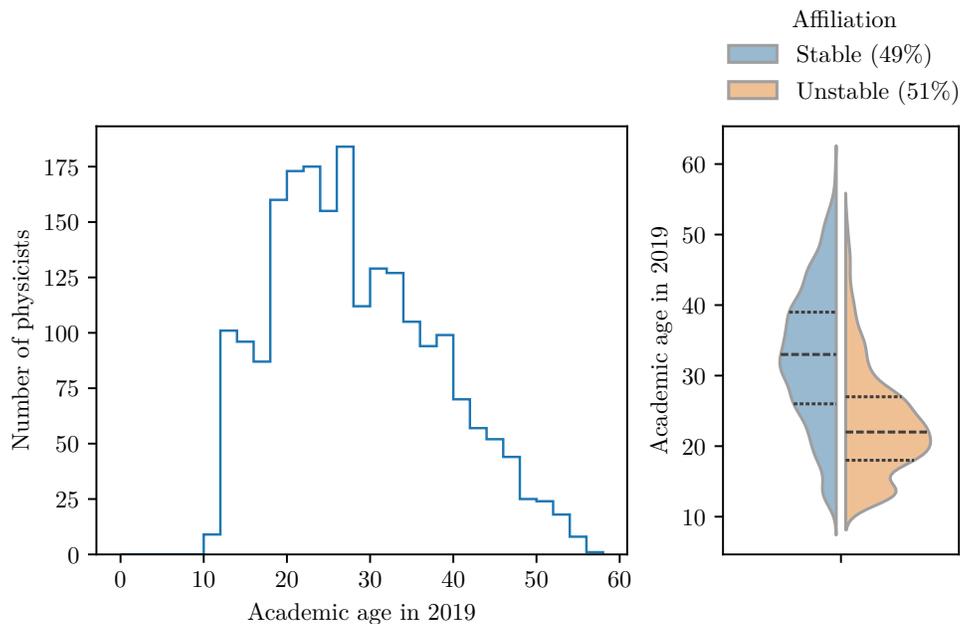


Figure 11: **Distribution of academic age (in 2019) and affiliation stability within the considered cohort.** Academic age (left plot) is evaluated as the time passed since the first publication. There is a cut-off in the distribution academic age due to the requirement on publication counts between 2000 and 2009. The median academic age is 27 years in 2019. 49% of the cohort has had at least one permanent affiliation spanning the period under consideration (right plot). As expected, permanent affiliations are associated with higher academic ages (bars indicate quartiles).

A.3 Topics

A.3.1 Word embeddings dimension

The word2vec skip-gram models is trained using different values for L the dimension of the embeddings' space. $L = 50$ (the choice made in the present paper) generally lies somewhere between under-fitting and over-fitting. The latter is a concern due to the small sample size (the model is trained on abstracts rather than full-texts).

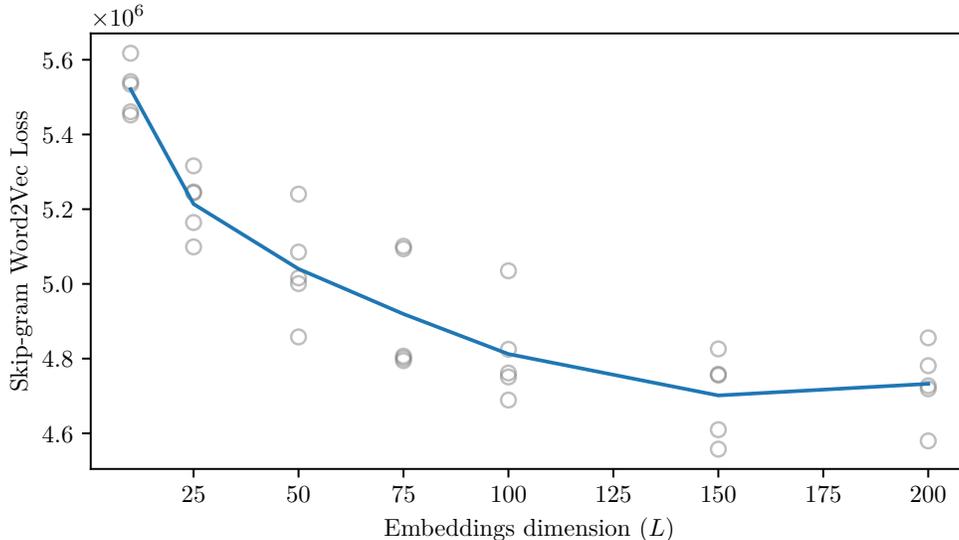


Figure 12: Loss of the word2vec model as a function of the embeddings dimension (L).

A.3.2 Keyword classification

In the Embedding Topic Model used in the present paper – as in topic models in general –, documents are mixtures of several topics; moreover, keywords may belong to several topics. This is a desirable feature: certain concepts serve different purposes depending on the context, and some concepts do not clearly belong to any research area. More importantly, this is crucial to the present approach, given that we must be sensitive to situations in which scientists have repurposed certain concepts to new research goals; i.e., instances where the same resources are applied in a new context (a new research area). For that, each keyword i from a document d is assigned a research area $z_{di} := \arg \max_{k=1, \dots, K} P(z_{di} = k | w_{di}, \boldsymbol{\theta}_d)$, which is the research area most probably associated with the keyword i given the topic distribution of the document $\boldsymbol{\theta}_d$ – i.e. the context. In fact:

$$z_{di} := \arg \max_{k=1, \dots, K} P(z_{di} = k | w_{di}, \boldsymbol{\theta}_d) = \arg \max_{k=1, \dots, K} \frac{P(w_{di} | z_{di} = k) P(z_{di} = k | \boldsymbol{\theta}_d)}{P(w_{di} | \boldsymbol{\theta}_d)} \quad (14)$$

In the process, we discard ambiguous keywords for which $H(z_{di}) \geq \log 2$, where H denotes the entropy of the distribution $P(z_{di} = k | w_{di})$. Either such keywords do not carry any scientific content, or the context is insufficient to disambiguate among the possible research areas to which they might belong.

The average effective amount of topics in documents according to the topic model – measured as $\exp H(\boldsymbol{\theta}_d)$ where H is the Shannon entropy – is 7.5, which is unrealistically high. The filtering and classification procedure reduces the average effective amount of topics to 3.1 per document, which is much more informative. This procedure is especially important for short texts such as abstracts, which poorly constrain the latent topic distribution $\boldsymbol{\theta}_d$.

A.3.3 List of topics

Table 1: Research areas, their top-words, and their correlation with a standard classification (PACS).

Research area	Top words	Most correlated PACS categories
AdS/CFT	boundary, holographic, flow, bulk, critical, conformal, critical_point, boundary_theory, cfts, point, bootstrap, central, conformal_anomaly, strip, fixed, free, entanglement_entropy, conformal_field_theory, criticality, condition	Gauge/string duality (0.27) Conformal field theory, algebraic [...] (0.23) Theory of quantized fields (0.15) Theories and models of [...] (0.14) Critical point phenomena (0.14)
Black holes	hole, black_hole, gravity, black, horizon, geometry, gravitational, spacetimes, spacetime, curvature, thermodynamics, einstein, schwarzschild, metric, ad, hair, relativity, space_time, observer, graviton	Quantum aspects of black holes, [...] (0.51) Classical black holes (0.41) Physics of black holes (0.32) Exact solutions (0.24) Higher-dimensional black holes, [...] (0.23)
Classical fields	scalar, general, first, scalar_field, massless, real, explicit, dynamical, second, exact, special, linear, full, symmetric, static, electromagnetic, classical, nonlinear, approximate, non_trivial	Modified theories of gravity (0.19) Lorentz and Poincaré invariance (0.16) Nonlinear or nonlocal theories and [...] (0.11) Exact solutions (0.10) Higher-dimensional gravity and [...] (0.10)
Collider physics	distribution, collision, production, cross_sections, section, parton, hadron, cross_section, process, hadronic_collision, scattering, correction, fragmentation, partons, kinematics, transverse, impact, event, partonic	Perturbative calculations (0.29) Polarization in interactions and [...] (0.28) Inclusive production with [...] (0.27) Total and inclusive cross sections [...] (0.26) Relativistic heavy-ion collisions (0.25)
Cosmology	constant, cosmological, inflation, cosmic, perturbation, vacuum, universe, inflationary, cosmology, fluctuation, inhomogeneity, tension, lambda, planck, inflaton, cosmological_perturbation, era, epoch, density, background	Particle-theory and field-theory [...] (0.59) Cosmology (0.32) Observational cosmology (including [...] (0.28) Dark energy (0.25) Background radiations (0.21)
Dark matter	matter, dark, dark_matter, detection, dm, signal, abundance, observation, relic, direct, constraint, candidate, wimp, asymmetric, prospect, dark_matter_particle, center, detectable, cold, contribute	Dark matter (0.74) -ray (0.22) Cosmic rays (0.19) -ray sources; -ray bursts (0.17) Elementary particle processes (0.17)
Electroweak sector	standard, higgs, boson, particle, standard_model, physic, lhc, sm, top, tev, collider, mssm, electroweak, minimal, phenomenology, extension, extra, supersymmetric_model, superpartners, new_particle	Extensions of electroweak Higgs sector (0.34) Supersymmetric models (0.33) Non-standard-model Higgs bosons (0.30) Supersymmetric partners of known [...] (0.28) Standard-model Higgs bosons (0.27)
Gauge theory & Grand Unification	dimension, coupling, scale, structure, operator, fermion, value, matrix, number, su, charge, sector, spin, group, topological, anomalous, breaking, anomaly, global, flavor	Unified theories and models of [...] (0.22) Unification of couplings; mass relations (0.17) Quark and lepton masses and mixing (0.13) Unified field theories and models (0.12) Field theories in dimensions other [...] (0.12)
Hadrons	decay, data, channel, bound, resonance, gamma, meson, width, experimental_data, collaboration, kaon, prediction, experimental, measurement, admixture, narrow, process, hadronic_decay, s0, ratio	Decays of bottom mesons (0.30) Decays of J/, , and other quarkonia (0.24) Meson-meson interactions (0.21) Decays of bottom mesons (0.20) Bottom mesons ($ B > 0$) (0.19) Neutrino mass and mixing (0.74)
Neutrinos & flavour physics	neutrino, violation, oscillation, flavor, cp, angle, mixing, experiment, lepton, flavour, hierarchy, majorana, cp_violation, beta, leptogenesis, asymmetry, neutrino_mass, neutrino_oscillation, smallness, generation	Non-standard-model neutrinos, [...] (0.41) Ordinary neutrinos (0.30) Neutrino interactions (0.28) Quark and lepton masses and mixing (0.23)
Perturbative methods	amplitude, qcd, loop, diagram, sum, contribution, perturbative, expansion, vertex, rule, light_cone, perturbative_qcd, propagator, approach, correlator, one_loop, evaluation, non-perturbative, kernel, diagrammatic	General properties of perturbation [...] (0.25) Other nonperturbative calculations (0.24) Sum rules (0.22) Perturbative calculations (0.21) General properties of QCD [...] (0.16) Lattice QCD calculations (0.27)
QCD	quark, chiral, magnetic, baryon, relativistic, moment, qcd, light_quark, strong, heavy, heavy_quark, lattice, magnetic_field, electric, deconfinement, chromodynamics, current, diquarks, plasma, color	Chiral symmetries (0.26) Chiral Lagrangians (0.25) Quark-gluon plasma (0.23) General properties of QCD [...] (0.20)
Quantum Field Theory	quantum, group, quantum_field, representation, quantisation, mechanic, quantum_field_theory, transformation, hamiltonians, algebra, finite_dimensional, quantization, commutator, algebraic, arbitrary, operator, qft, invariant, analog, associated	Algebraic methods (0.26) Noncommutative field theory (0.25) Quantum mechanics (0.22) Noncommutative geometry (0.19) Quantum groups (0.18)

Continued on next page

Table 1: Research areas, their top-words, and their correlation with a standard classification (PACS).

Research area	Top words	Most correlated PACS categories
String theory & supergravity	string, supersymmetric, superstring, six_dimensional, modulus, super, instantons, supergravity, dyons, n2, mathcaln, superpotentials, heterotic, sigma_models, n1, n4, gauged, space, deformation, compactifications	Supersymmetry (0.31) Strings and branes (0.29) Supergravity (0.29) Compactification and four- [...] (0.25) D branes (0.20)
Thermodynamics	potential, effective, interaction, limit, temperature, action, finite, local, freedom, approximation, level, weak, chemical, force, effective_field_theory, lagrangian, finite_temperature, effective_field, degree, effective_theory	Finite-temperature field theory (0.26) Chiral symmetries (0.09) Nuclear forces (0.08) Quark-gluon plasma (0.08) General properties of QCD [...] (0.08) Lattice QCD calculations (0.07)
Uninterpretable	approach, method, analysis, recent, calculation, numerical, formalism, study, prediction, sigma, previous, work, theoretical, systematic, comparison, uncertainty, agreement, good, investigation, paper	Baryon resonances (S=C=B=0) (0.05) Other nonperturbative calculations (0.05) Few-body systems (0.05) Lagrangian and Hamiltonian approach (0.05)
Uninterpretable	solution, equation, phase, space, time, system, transition, region, condition, constraint, dynamic, class, background, configuration, wave, range, motion, set, star, instability	Exact solutions (0.14) Nonlinear or nonlocal theories and [...] (0.11) Extended classical solutions; [...] (0.10) Relativistic wave equations (0.10) Modified theories of gravity (0.09) Electromagnetic form factors (0.17)
Uninterpretable	form, correction, momentum, tensor, mode, relation, higher, factor, vector, invariant, formula, angular, part, theorem, spectrum, power, dimensional, invariance, expression, derivative	Protons and neutrons (0.10) Lorentz and Poincaré invariance (0.08) Gauge field theories (0.06) Dispersion relations (0.06)
Uninterpretable	spectrum, low, problem, low_energy, important, high, property, high_energy, small, physical, soft, fundamental, behavior, analytic, behaviour, spectral, dispersion, essential, phenomenon, regime	General properties of QCD [...] (0.07) Regge formalism (0.07) Wave propagation, transmission and [...] (0.05) Elastic scattering (0.05) Lattice gauge theory (0.05)
Uninterpretable	different, possible, particular, present, various, mechanism, type, example, massive, several, scenario, simple, single, similar, consistent, addition, hand, different_type, interesting, way	Particle-theory and field-theory [...] (0.07) Modified theories of gravity (0.07) Field theories in dimensions other [...] (0.05) Cosmology (0.05) Dark energy (0.05)

A.3.4 Topic validation using the citation network

In order to validate the consistency and scientific dimension of the topics that were recovered, we verify that papers from a given topic tend to cite more papers from the same topic. Let $N_{k,k'}$ be the amount of citations of articles that belong to topic k' originating from articles that belong to k , and $N = \sum_{k,k'} N_{k,k'}$ the total number of citations. From this matrix, a normalized pointwise mutual correlation $\text{npmi}(k, k')$ is calculated:

$$\text{npmi}(k, k') = \log \frac{N_{k,k'}/N}{(\sum_i N_{k,i}/N)(\sum_i N_{i,k'}/N)} \quad (15)$$

$\text{npmi}(k, k')$ is shown in Figure 13. It measures how frequent citations from k to k' are, relative to what would be expected if citations were uniformly distributed. The diagonal values are positive, indicating that the topics we retrieved tend to refer to themselves significantly more than expected by chance alone, providing further evidence of their scientific content and coherence.

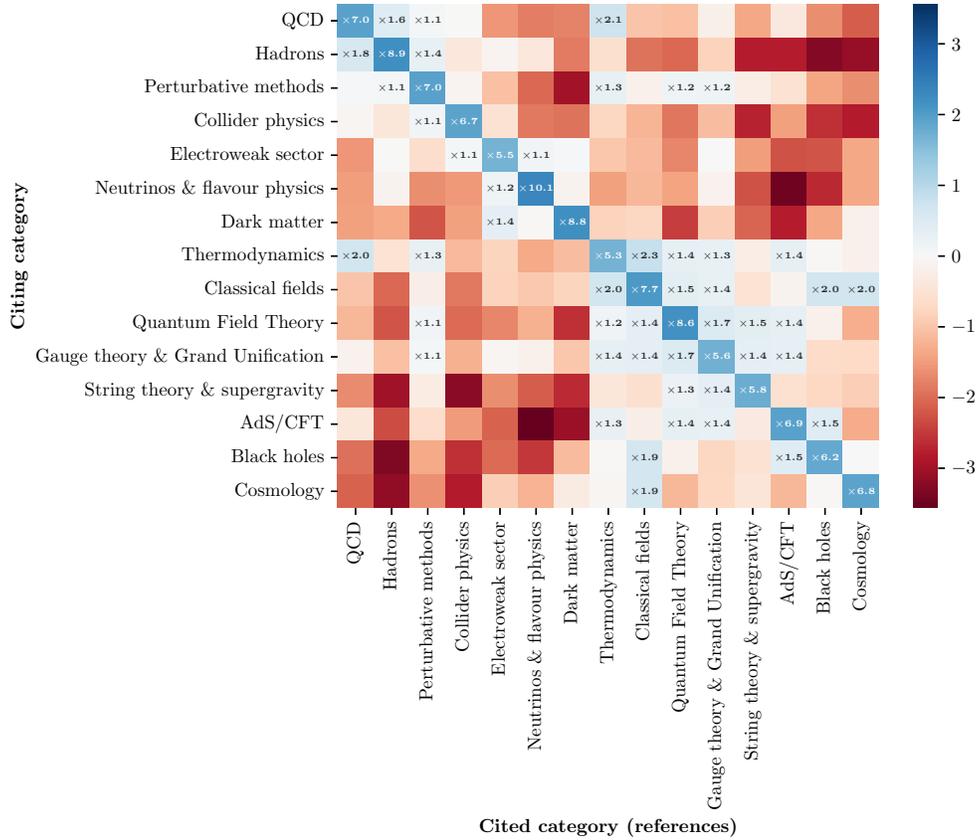


Figure 13: **Normalized pointwise mutual information $\text{npmi}(k, k')$ of the citation matrix $N_{k, k'}$.** Positive values (in blue) indicate that the y-axis research area cites the x-axis research area more often than expected by chance; negative values indicate that citations occur less than expected by chance. When shown, individual values indicate how many times citations occur compared to chance alone. For instance, papers about Hadrons cite papers about Perturbative methods 1.4 times what would be expected if citations were uniform across research areas.

A.3.5 A comparison of three topic models

Figure 14 compares the ability of three topic models to measure the transformations in high-energy physics research resulting from changes in the landscape of experimental opportunities. In the coarse-grained model ($K_0 = 15$, in the middle), many types of experiments are lumped together into a single topic. It is therefore ill-suited for assessing the impact of the transformations in the landscape of experimental opportunities. The model used throughout the paper is well able to distinguish neutrino research for dark matter research, which have both undergone significant transformations according to Figure 10. It is also better able to separate black hole phenomenology and cosmology, compared to the fine-grained model ($K_0 = 25$).

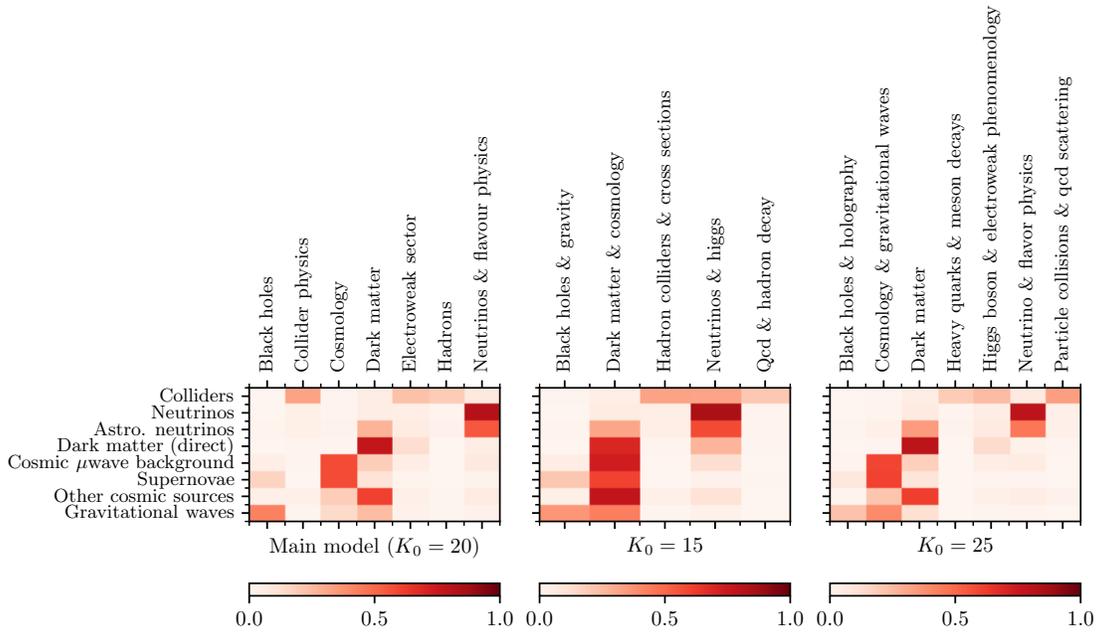


Figure 14: Origin of papers citing each type of experiment, according to three topic models: the model considered throughout the present paper (to the left), a model with fewer topics ($K_0 = 15$), and a model with more topics ($K_0 = 25$). A value of one (dark red) indicates that 100% of the papers citing a certain type of experiment (cf. rows) originate from a given topic (cf. columns).

A.3.6 Topic validation using the PACS classification



Figure 15: Correlation between PACS categories present in ≥ 100 publications in the corpus (rows), and the research areas recovered by the topic model (columns). Each colored cell indicates the correlation between a leaf category of the hierarchical PACS classification and one of the topics from the topic model.

Blue cells show that the research areas recovered from the topic model correlate with the PACS classification (which further confirms their scientific dimension), but also that they can give a different picture. For instance, each of the topics “dark matter”, “black holes” and “cosmology” span over several higher-level categories of the PACS classification (e.g. “fundamental astronomy ...” and “specific theories and interaction models ...” for dark matter).

A.4 Measures of capital

Figure 16 shows the Pearson correlation between different measures of the diversity of intellectual and social capital and of power, as evaluated among the cohort of high-energy physicists. For comparison purposes, an alternative measure of diversity based on Stirling’s index [68], with prior applications to studies of interdisciplinarity [69, 70] is evaluated²³. A measure of brokerage is also considered²⁴.

As shown in Figure 16, the entropic measure of diversity considered in this paper correlates strongly with the Stirling measure. The magnitude of social capital (which is similar to degree centrality) correlates weakly with excess social diversity, thus emphasizing that power and diversity are partially orthogonal aspects of social capital. The magnitude of social capital is strongly correlated with brokerage; indeed, strongly connected scientists, with higher degree centrality, are also those scientists who initiate collaborations between otherwise disconnected scientists, as measured by brokerage.

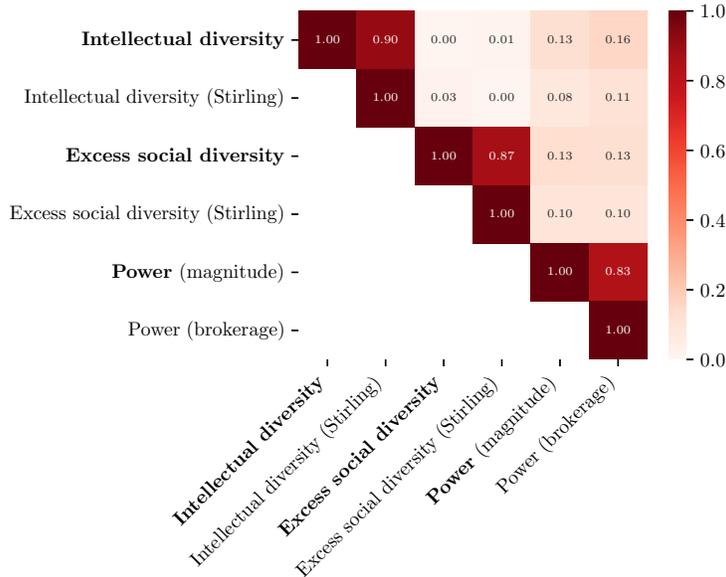


Figure 16: **Correlation between different measures of capital.** Measures considered in priority in this analysis are shown in bold. Alternative measures are shown for comparison purposes. By construction, excess social diversity is orthogonal to intellectual diversity ($R = 0$).

²³The Stirling-based diversity measured is evaluated as:

$$D_{\text{Stirling}} = 1 - \sum_{k,k'} d_{kk'} I_{ak} I_{ak'} \tag{16}$$

Where $d_{kk'}$ is the fraction of scientists who have more expertise than average in both k and k' among those that have expertise in one or the other (i.e., a similarity matrix). This follows from previous approaches for measuring research interdisciplinarity [69, 70].

²⁴We evaluated brokerage as the amount of pairs of scientists that have collaborated with a given physicist while having no common collaborator except for this physicist. This effectively measures the extent to which this physicist connects otherwise disconnected scientists.

A.5 Model performance over multiple corpora, temporal segmentations, and topic granularities

The predictive power of the model can be assessed by evaluating the total variation distance between the true distribution \mathbf{y}_a and the predicted distribution $\mathbf{y}_a^{\text{pred}}$. This performance metric is calculated via 10-fold cross-validation. It is compared to a baseline model that predicts no change in the research agenda ($\mathbf{y}_a^{\text{baseline}} = \mathbf{x}_a$). The results are shown in Table 2. For the cohort of high-energy physicists, the model performs only marginally better than the baseline, given that individuals have remained quite conservative on average, most of the fluctuations being difficult to predict. Table 2 also considers a cohort of scientists from the ACL anthology corpus of computational linguistics research [71], by running the same pipeline (the measurement of research portfolios during two consecutive time periods using our topic model approach and the training of the trajectory model). Although the data are of significantly lesser quality, research portfolios have undergone much more significant transformations in this dataset (see Appendix A.10, Figure 20). Consequently, our model performs much better than the baseline for this cohort.

Table 2: Performance of the actual model versus that of the baseline model for various corpora, temporal segmentations, topic model parameters, and authorship criteria.

Corpus	Authorship	K	L	Cohort size	Temporal segmentation	Model $\mu(d_{\text{TV}}(\mathbf{y}_a, \mathbf{y}_a^{\text{pred}}))$	Baseline $\mu(d_{\text{TV}}(\mathbf{y}_a, \mathbf{x}_a))$
HEP	Any	20	50	2108	2000–2009 2015–2019	0.306	0.316
HEP	1st/last	20	50	2108	2000–2009 2015–2019	0.306	0.316
HEP	Any	20	50	1836	2000–2004 2005–2009	0.262	0.262
HEP	Any	20	50	2530	2005–2009 2010–2014	0.261	0.265
HEP	Any	20	50	3816	2010–2014 2015–2019	0.246	0.244
HEP	Any	15	50	2375	2000–2009 2015–2019	0.293	0.297
HEP	Any	25	50	2109	2000–2009 2010–2019	0.315	0.328
HEP	Any	15	50	2069	2000–2009 2015–2019	0.290	0.295
HEP	Any	20	150	2169	2000–2009 2015–2019	0.309	0.318
ACL Anthology	Any	20	50	578	2002–2011 2012–2022	0.337	0.466

A.6 Learning costs and optimal transport

The MCMC algorithm from [17] is run on 1 000 000 iterations of the “MetroMC” algorithm, using what the authors call a “P1” prior (that is, a prior such that $\sum_{kk'} C_{kk'} = C_0 = \text{cst}^{25}$). More precisely, we assume that:

²⁵We chose the minimum value of C_0 for which the system admitted a solution.

$$P(c_{kk'}|p_{kk'}) = \frac{1}{Z} \frac{1}{\prod_{kk'} c_{kk'}^{1/2}} \exp(-\alpha D_{KL}(c_{kk'}||p_{kk'})) \text{ with } c_{kk'} = C_{kk'}/C_0 \quad (17)$$

$$\text{and } p_{kk'} = \text{softmax}(\beta(1 - \nu_{kk'})) \quad (18)$$

$$(19)$$

(17) is sometimes referred to as the entropic prior [72, 73]. The mean posterior values of $C_{k,k'}$ are shown in Figure 17, as a function of the knowledge gap from k to k' . The knowledge gap $1 - \nu_{kk'}$ is the fraction of experts in k who do not hold significant expertise in k' (ν is shown in Figure 6b). A significant correlation is found ($R = -0.76$). This is true also for the replication dataset of computational linguistics research (Appendix A.10), for which $R = -0.63$.

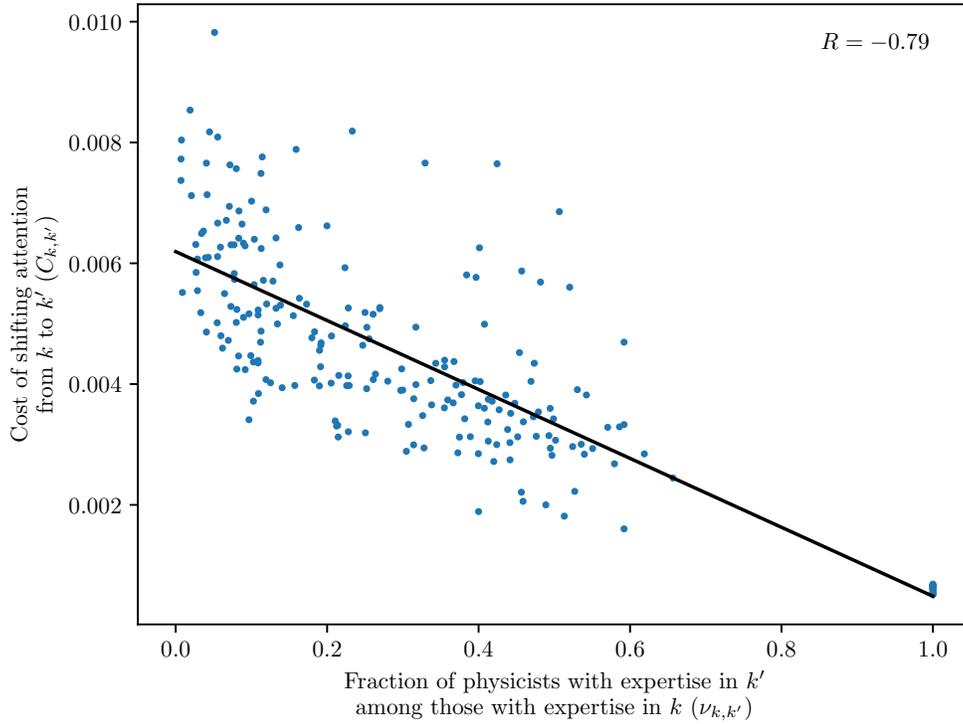


Figure 17: Cost of shifting a unit of cohort’s research attention from k to k' as a function of the fraction of physicists with expertise in k' among those with expertise in k ($\nu_{k,k'}$).

A.7 Effect of capital on strategies of change

A.7.1 Model for the magnitude of change

The model for c_a is:

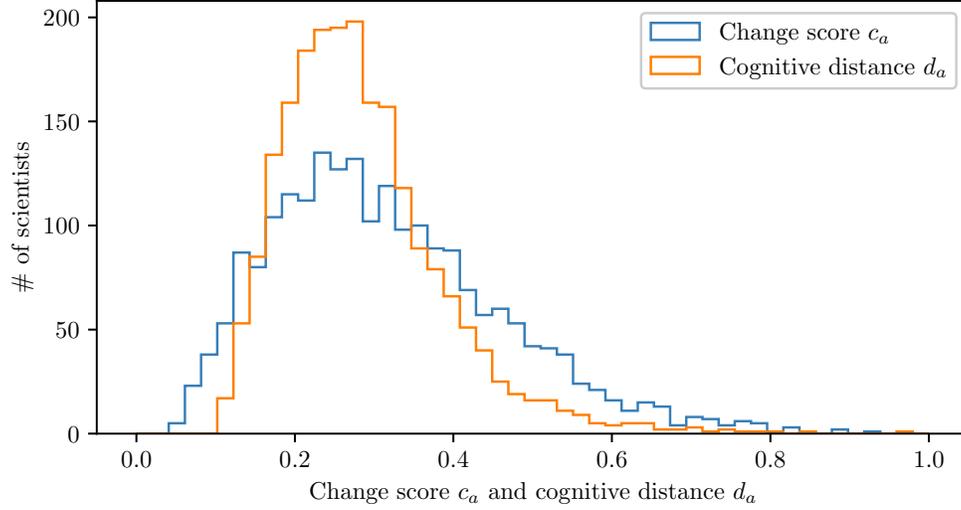


Figure 18: **Distribution of change and cognitive distances in the cohort.** Higher values correspond to more drastic changes in a scientists' research agenda.

$$z(c_a) \sim \mathcal{N}(\mu_a, \sigma)$$

$$\mu_a = \beta^{\text{int-div}} z(D(\mathbf{I}_a)) + \beta^{\text{soc-div}} z(D^*(\mathbf{S}_a)) + \beta^{\text{power}} z(P(\mathbf{S}_a)) + \beta^{\text{stability}} p_a + \beta^{\text{age}} z(a_a) + \beta^{\text{prod}} z(\pi_a) + \mu_{k_a}^{\text{area}} + \mu$$

$$k_a = \arg \max_k x_{ak}$$

$$\beta, \mu \sim \mathcal{N}(0, 1)$$

$$|\mu_k^{\text{area}}| \sim \text{Exponential}(\tau)$$

$$\tau, \sigma \sim \text{Exponential}(1)$$

Where $z(\cdot)$ denotes standardized variables.

A.7.2 Model for the probability of having entered/exited a research area

The model for the probability p_a of having entered a new field:

$$\text{logit}(p_a) = \beta^{\text{int-div}} z(D(\mathbf{I}_a)) + \beta^{\text{soc-div}} z(D^*(\mathbf{S}_a)) + \beta^{\text{power}} z(P(\mathbf{S}_a)) + \beta^{\text{stability}} p_a + \beta^{\text{age}} z(a_a) + \beta^{\text{prod}} z(\pi_a) + \mu_{k_a}^{\text{area}} + \mu$$

$$k_a = \arg \max_k x_{ak}$$

$$\beta, \mu \sim \mathcal{N}(0, 1)$$

$$|\mu_k^{\text{area}}| \sim \text{Exponential}(\tau)$$

$$\tau, \sigma \sim \text{Exponential}(1)$$

The same model structure is used for the probability of having exited a research area.

A.7.3 Effect of capital and robustness checks

Table 3: Effect of each variable on (a) the change score and (b) the cognitive distance for each model. The reference model uses entropy as the diversity measure D and the magnitude of intellectual capital as a measure of power P . Values indicate the mean posterior effect size and the 95% credible interval. Significant effects are shown in bold.

Dep. variable	Change score (c_a)			Cognitive distance (d_a)		
Model	Reference	$D = \text{Stirling}$	$P = \text{Brokerage}$	Reference	$D = \text{Stirling}$	$P = \text{Brokerage}$
Predictor						
Intellectual capital (diversity)	+0.28 ^{+0.044} _{-0.044}	+0.28 ^{+0.042} _{-0.043}	+0.27 ^{+0.044} _{-0.043}	+0.33 ^{+0.043} _{-0.042}	+0.34 ^{+0.042} _{-0.042}	+0.32 ^{+0.043} _{-0.043}
Social capital (diversity)	+0.09 ^{+0.04} _{-0.04}	+0.07 ^{+0.04} _{-0.04}	+0.08 ^{+0.04} _{-0.04}	+0.11 ^{+0.04} _{-0.041}	+0.09 ^{+0.04} _{-0.04}	+0.1 ^{+0.04} _{-0.04}
Social capital (power)	-0.09 ^{+0.06} _{-0.06}	-0.07 ^{+0.06} _{-0.06}	-0.02 ^{+0.05} _{-0.05}	-0.14 ^{+0.061} _{-0.061}	-0.12 ^{+0.06} _{-0.06}	-0.05 ^{+0.05} _{-0.05}
Stable affiliation	-0.01 ^{+0.09} _{-0.09}	-0.009 ^{+0.09} _{-0.09}	-0.0008 ^{+0.09} _{-0.09}	-0.007 ^{+0.09} _{-0.09}	+0.0009 ^{+0.09} _{-0.09}	+0.01 ^{+0.09} _{-0.09}
Academic age	-0.1 ^{+0.05} _{-0.05}	-0.1 ^{+0.05} _{-0.05}	-0.1 ^{+0.047} _{-0.047}	-0.07 ^{+0.05} _{-0.05}	-0.07 ^{+0.047} _{-0.047}	-0.08 ^{+0.05} _{-0.05}
Productivity (co-authored)	-0.12 ^{+0.058} _{-0.059}	-0.12 ^{+0.058} _{-0.058}	-0.17 ^{+0.052} _{-0.053}	-0.1 ^{+0.06} _{-0.06}	-0.1 ^{+0.058} _{-0.056}	-0.17 ^{+0.053} _{-0.052}
Productivity (solo-authored)	-0.05 ^{+0.041} _{-0.04}	-0.05 ^{+0.04} _{-0.04}	-0.06 ^{+0.04} _{-0.04}	-0.04 ^{+0.04} _{-0.04}	-0.03 ^{+0.04} _{-0.04}	-0.04 ^{+0.04} _{-0.04}
Hadrons	-0.009 ^{+0.2} _{-0.2}	-0.11 ^{+0.18} _{-0.2}	-0.008 ^{+0.2} _{-0.2}	+0.03 ^{+0.1} _{-0.1}	-0.05 ^{+0.1} _{-0.2}	+0.04 ^{+0.1} _{-0.1}
String theory & supergravity	+0.28 ^{+0.15} _{-0.15}	+0.34 ^{+0.18} _{-0.18}	+0.25 ^{+0.15} _{-0.15}	+0.11 ^{+0.13} _{-0.11}	+0.21 ^{+0.15} _{-0.15}	+0.07 ^{+0.1} _{-0.1}
Perturbative methods	+0.12 ^{+0.22} _{-0.18}	+0.06 ^{+0.2} _{-0.2}	+0.13 ^{+0.22} _{-0.18}	+0.15 ^{+0.22} _{-0.17}	+0.1 ^{+0.21} _{-0.17}	+0.16 ^{+0.23} _{-0.18}
Classical fields	-0.25 ^{+0.36} _{-0.59}	-0.21 ^{+0.38} _{-0.55}	-0.23 ^{+0.35} _{-0.57}	-0.19 ^{+0.27} _{-0.58}	-0.17 ^{+0.3} _{-0.51}	-0.16 ^{+0.26} _{-0.55}
Collider physics	-0.19 ^{+0.16} _{-0.17}	-0.34 ^{+0.19} _{-0.19}	-0.2 ^{+0.17} _{-0.17}	-0.02 ^{+0.1} _{-0.1}	-0.16 ^{+0.15} _{-0.16}	-0.03 ^{+0.1} _{-0.1}
Neutrinos & flavour physics	+0.21 ^{+0.18} _{-0.17}	+0.17 ^{+0.2} _{-0.18}	+0.18 ^{+0.18} _{-0.17}	+0.11 ^{+0.16} _{-0.13}	+0.1 ^{+0.2} _{-0.1}	+0.09 ^{+0.2} _{-0.1}
Black holes	+0.06 ^{+0.2} _{-0.2}	+0.15 ^{+0.22} _{-0.19}	+0.05 ^{+0.2} _{-0.2}	-0.03 ^{+0.1} _{-0.2}	+0.06 ^{+0.2} _{-0.1}	-0.04 ^{+0.1} _{-0.2}
Gauge theory & Grand Unification	-0.07 ^{+0.3} _{-0.4}	-0.05 ^{+0.3} _{-0.4}	-0.06 ^{+0.3} _{-0.3}	-0.02 ^{+0.2} _{-0.3}	+0.003 ^{+0.3} _{-0.3}	-0.008 ^{+0.2} _{-0.3}
Dark matter	-0.27 ^{+0.24} _{-0.25}	-0.32 ^{+0.26} _{-0.27}	-0.28 ^{+0.24} _{-0.25}	-0.13 ^{+0.17} _{-0.23}	-0.18 ^{+0.2} _{-0.24}	-0.13 ^{+0.17} _{-0.23}
Thermodynamics	+0.14 ^{+0.36} _{-0.26}	+0.27 ^{+0.41} _{-0.32}	+0.16 ^{+0.37} _{-0.27}	+0.1 ^{+0.34} _{-0.2}	+0.25 ^{+0.41} _{-0.29}	+0.12 ^{+0.36} _{-0.21}
Cosmology	-0.02 ^{+0.16} _{-0.17}	+0.01 ^{+0.2} _{-0.2}	-0.03 ^{+0.2} _{-0.2}	-0.06 ^{+0.1} _{-0.2}	-0.02 ^{+0.1} _{-0.2}	-0.06 ^{+0.13} _{-0.17}
Electroweak sector	-0.14 ^{+0.15} _{-0.16}	-0.21 ^{+0.18} _{-0.18}	-0.17 ^{+0.16} _{-0.16}	-0.04 ^{+0.1} _{-0.1}	-0.1 ^{+0.13} _{-0.15}	-0.07 ^{+0.1} _{-0.1}
QCD	-0.009 ^{+0.2} _{-0.2}	-0.07 ^{+0.2} _{-0.2}	+0.001 ^{+0.2} _{-0.2}	+0.02 ^{+0.1} _{-0.1}	-0.03 ^{+0.14} _{-0.16}	+0.04 ^{+0.1} _{-0.1}
Quantum Field Theory	+0.04 ^{+0.2} _{-0.2}	+0.09 ^{+0.2} _{-0.2}	+0.04 ^{+0.21} _{-0.19}	-0.09 ^{+0.2} _{-0.2}	-0.04 ^{+0.2} _{-0.2}	-0.08 ^{+0.2} _{-0.2}
AdS/CFT	+0.003 ^{+0.2} _{-0.2}	+0.07 ^{+0.3} _{-0.2}	+0.005 ^{+0.2} _{-0.2}	+0.02 ^{+0.2} _{-0.2}	+0.11 ^{+0.27} _{-0.2}	+0.03 ^{+0.2} _{-0.2}

Table 4: Effect of each variable on (a) the probability of having entered a new research area and (b) the probability of having exited a research area, for each model. The reference model uses entropy as the diversity measure D and the magnitude of intellectual capital as a measure of power P . Values indicate the mean posterior effect size and the 95% credible interval. Significant effects are shown in bold.

Dep. variable Model	Entered a new research area			Exited a research area		
	Reference	$D = \text{Stirling}$	$P = \text{Brokerage}$	Reference	$D = \text{Stirling}$	$P = \text{Brokerage}$
Predictor						
Intellectual capital (diversity)	+0.2 ^{+0.11} _{-0.11}	+0.17 ^{+0.1} _{-0.1}	+0.19 ^{+0.11} _{-0.11}	+1 ^{+0.14} _{-0.14}	+0.85 ^{+0.12} _{-0.12}	+1 ^{+0.14} _{-0.14}
Social capital (diversity)	+0.22 ^{+0.1} _{-0.1}	+0.18 ^{+0.1} _{-0.1}	+0.22 ^{+0.099} _{-0.1}	+0.04 ^{+0.1} _{-0.1}	+0.04 ^{+0.1} _{-0.1}	+0.04 ^{+0.1} _{-0.1}
Social capital (power)	+0.006 ^{+0.1} _{-0.1}	+0.03 ^{+0.15} _{-0.15}	+0.04 ^{+0.1} _{-0.1}	-0.03 ^{+0.2} _{-0.2}	+0.02 ^{+0.2} _{-0.2}	+0.03 ^{+0.1} _{-0.1}
Stable affiliation	-0.19 ^{+0.22} _{-0.22}	-0.18 ^{+0.22} _{-0.22}	-0.19 ^{+0.22} _{-0.22}	+0.04 ^{+0.2} _{-0.2}	+0.06 ^{+0.2} _{-0.2}	+0.04 ^{+0.24} _{-0.24}
Academic age	+0.04 ^{+0.12} _{-0.11}	+0.04 ^{+0.1} _{-0.1}	+0.04 ^{+0.1} _{-0.1}	-0.21 ^{+0.12} _{-0.12}	-0.21 ^{+0.12} _{-0.12}	-0.22 ^{+0.12} _{-0.12}
Productivity (co-authored)	-0.07 ^{+0.1} _{-0.1}	-0.08 ^{+0.1} _{-0.1}	-0.09 ^{+0.1} _{-0.1}	-0.28 ^{+0.15} _{-0.14}	-0.28 ^{+0.15} _{-0.15}	-0.31 ^{+0.13} _{-0.13}
Productivity (solo-authored)	-0.05 ^{+0.1} _{-0.09}	-0.05 ^{+0.1} _{-0.09}	-0.06 ^{+0.1} _{-0.1}	-0.02 ^{+0.1} _{-0.1}	-0.007 ^{+0.1} _{-0.1}	-0.03 ^{+0.1} _{-0.1}
Hadrons	-0.14 ^{+0.29} _{-0.36}	-0.24 ^{+0.33} _{-0.39}	-0.14 ^{+0.29} _{-0.35}	+0.03 ^{+0.3} _{-0.2}	-0.09 ^{+0.3} _{-0.4}	+0.04 ^{+0.3} _{-0.2}
String theory & supergravity	+0.32 ^{+0.32} _{-0.3}	+0.4 ^{+0.35} _{-0.33}	+0.32 ^{+0.32} _{-0.29}	+0.3 ^{+0.34} _{-0.3}	+0.65 ^{+0.39} _{-0.38}	+0.28 ^{+0.34} _{-0.29}
Perturbative methods	+0.11 ^{+0.45} _{-0.35}	+0.09 ^{+0.5} _{-0.4}	+0.11 ^{+0.45} _{-0.35}	-0.03 ^{+0.29} _{-0.34}	-0.13 ^{+0.4} _{-0.48}	-0.03 ^{+0.3} _{-0.3}
Classical fields	+0.22 ^{+1.1} _{-0.6}	+0.34 ^{+1.4} _{-0.7}	+0.22 ^{+1.1} _{-0.6}	-0.07 ^{+0.4} _{-0.6}	-0.07 ^{+0.7} _{-0.8}	-0.07 ^{+0.4} _{-0.6}
Collider physics	-0.43 ^{+0.33} _{-0.34}	-0.61 ^{+0.33} _{-0.34}	-0.42 ^{+0.33} _{-0.34}	-0.01 ^{+0.2} _{-0.2}	-0.28 ^{+0.32} _{-0.37}	-0.02 ^{+0.2} _{-0.2}
Neutrinos & flavour physics	+0.08 ^{+0.3} _{-0.3}	+0.04 ^{+0.3} _{-0.3}	+0.07 ^{+0.3} _{-0.3}	-0.21 ^{+0.25} _{-0.35}	-0.31 ^{+0.35} _{-0.41}	-0.21 ^{+0.26} _{-0.36}
Black holes	-0.0006 ^{+0.3} _{-0.3}	+0.06 ^{+0.4} _{-0.3}	-0.003 ^{+0.3} _{-0.3}	+0.08 ^{+0.4} _{-0.3}	+0.43 ^{+0.53} _{-0.45}	+0.08 ^{+0.4} _{-0.3}
Gauge theory & Grand Unification	-0.04 ^{+0.6} _{-0.6}	-0.04 ^{+0.6} _{-0.7}	-0.03 ^{+0.6} _{-0.6}	-0.08 ^{+0.4} _{-0.6}	-0.11 ^{+0.64} _{-0.79}	-0.08 ^{+0.4} _{-0.6}
Dark matter	-0.62 ^{+0.55} _{-0.56}	-0.68 ^{+0.55} _{-0.56}	-0.63 ^{+0.55} _{-0.56}	-0.05 ^{+0.3} _{-0.4}	-0.11 ^{+0.42} _{-0.5}	-0.05 ^{+0.3} _{-0.4}
Thermodynamics	-0.03 ^{+0.5} _{-0.6}	+0.01 ^{+0.62} _{-0.58}	-0.02 ^{+0.5} _{-0.6}	-0.05 ^{+0.4} _{-0.6}	+0.03 ^{+0.7} _{-0.6}	-0.05 ^{+0.41} _{-0.53}
Cosmology	-0.07 ^{+0.3} _{-0.4}	-0.03 ^{+0.3} _{-0.4}	-0.07 ^{+0.3} _{-0.4}	+0.09 ^{+0.4} _{-0.3}	+0.36 ^{+0.56} _{-0.43}	+0.09 ^{+0.4} _{-0.3}
Electroweak sector	+0.05 ^{+0.3} _{-0.3}	+0.009 ^{+0.3} _{-0.3}	+0.05 ^{+0.3} _{-0.3}	-0.003 ^{+0.2} _{-0.2}	-0.04 ^{+0.3} _{-0.3}	-0.009 ^{+0.2} _{-0.2}
QCD	+0.008 ^{+0.3} _{-0.3}	-0.03 ^{+0.3} _{-0.4}	+0.01 ^{+0.3} _{-0.3}	+0.04 ^{+0.32} _{-0.26}	+0.04 ^{+0.4} _{-0.4}	+0.04 ^{+0.3} _{-0.3}
Quantum Field Theory	+0.15 ^{+0.52} _{-0.38}	+0.22 ^{+0.57} _{-0.41}	+0.15 ^{+0.52} _{-0.38}	+0.05 ^{+0.4} _{-0.3}	+0.26 ^{+0.64} _{-0.44}	+0.05 ^{+0.4} _{-0.3}
AdS/CFT	+0.14 ^{+0.6} _{-0.42}	+0.22 ^{+0.67} _{-0.47}	+0.14 ^{+0.59} _{-0.42}	-0.06 ^{+0.4} _{-0.5}	+0.07 ^{+0.6} _{-0.5}	-0.06 ^{+0.3} _{-0.5}

A.8 Additional robustness checks

The robustness of the results of the comparative analysis is assessed by varying different parameters:

- The papers included in each authors' portfolio (any paper versus first-authored and last-authored papers only).
- The amount of topics in the topic model (K_0).
- The amount of dimensions for the word embeddings (L).
- The temporal segmentation for the early and late research portfolios.

Table 5: Summary of the effect of each predictor on the change score (c_a) across topic models and temporal segmentations. Values indicate the mean posterior effect size and the 95% credible interval. Significant effects are shown in bold.

Author-ship	K_0	L	Temporal segmentation	Intell. capital (diversity)	Soc. capital (diversity)	Soc. capital (power)	Stable affiliation	Academic age	Prod. (co-auth.)	Prod. (solo-auth.)
Any	20	50	2000-2009 2015-2019	+0.28 ^{+0.044} _{-0.044}	+0.09 ^{+0.04} _{-0.04}	-0.09 ^{+0.06} _{-0.06}	-0.01 ^{+0.09} _{-0.09}	-0.1 ^{+0.05} _{-0.05}	-0.12 ^{+0.058} _{-0.059}	-0.05 ^{+0.041} _{-0.04}
1st/last	20	50	2000-2009 2015-2019	+0.25 ^{+0.055} _{-0.056}	+0.09 ^{+0.05} _{-0.05}	-0.01 ^{+0.08} _{-0.08}	+0.04 ^{+0.1} _{-0.1}	-0.12 ^{+0.059} _{-0.06}	-0.16 ^{+0.075} _{-0.074}	-0.05 ^{+0.05} _{-0.05}
Any	20	50	2000-2004 2005-2009	+0.37 ^{+0.045} _{-0.045}	+0.12 ^{+0.043} _{-0.043}	-0.02 ^{+0.06} _{-0.06}	-0.02 ^{+0.1} _{-0.1}	-0.08 ^{+0.05} _{-0.05}	-0.21 ^{+0.059} _{-0.059}	-0.06 ^{+0.04} _{-0.04}
Any	20	50	2010-2014 2015-2019	+0.36 ^{+0.039} _{-0.039}	+0.08 ^{+0.04} _{-0.04}	-0.06 ^{+0.051} _{-0.052}	-0.11 ^{+0.087} _{-0.086}	-0.06 ^{+0.04} _{-0.04}	-0.21 ^{+0.051} _{-0.051}	-0.04 ^{+0.04} _{-0.04}
Any	20	50	2000-2009 2010-2019	+0.37 ^{+0.033} _{-0.033}	+0.06 ^{+0.03} _{-0.03}	-0.06 ^{+0.05} _{-0.05}	-0.02 ^{+0.07} _{-0.07}	-0.04 ^{+0.036} _{-0.035}	-0.22 ^{+0.046} _{-0.046}	-0.04 ^{+0.03} _{-0.03}
Any	20	50	2000-2009 2015-2019	+0.32 ^{+0.044} _{-0.044}	+0.09 ^{+0.04} _{-0.04}	-0.09 ^{+0.06} _{-0.06}	+0.02 ^{+0.09} _{-0.09}	-0.12 ^{+0.047} _{-0.046}	-0.15 ^{+0.057} _{-0.057}	-0.05 ^{+0.04} _{-0.04}
Any	15	50	2000-2009 2015-2019	+0.33 ^{+0.04} _{-0.04}	+0.09 ^{+0.04} _{-0.04}	-0.12 ^{+0.057} _{-0.057}	-0.04 ^{+0.09} _{-0.09}	-0.12 ^{+0.05} _{-0.051}	-0.1 ^{+0.054} _{-0.054}	-0.06 ^{+0.04} _{-0.04}
Any	25	50	2000-2009 2015-2019	+0.35 ^{+0.042} _{-0.042}	+0.13 ^{+0.04} _{-0.04}	-0.15 ^{+0.058} _{-0.058}	-0.002 ^{+0.09} _{-0.09}	-0.15 ^{+0.052} _{-0.052}	-0.12 ^{+0.056} _{-0.056}	-0.05 ^{+0.04} _{-0.04}
Any	15	150	2000-2009 2015-2019	+0.35 ^{+0.044} _{-0.044}	+0.07 ^{+0.04} _{-0.04}	-0.08 ^{+0.06} _{-0.06}	-0.02 ^{+0.089} _{-0.089}	-0.08 ^{+0.05} _{-0.05}	-0.1 ^{+0.056} _{-0.057}	-0.02 ^{+0.04} _{-0.04}
Any	20	150	2000-2009 2015-2019	+0.34 ^{+0.044} _{-0.043}	+0.08 ^{+0.04} _{-0.04}	-0.09 ^{+0.06} _{-0.06}	+0.008 ^{+0.09} _{-0.09}	-0.1 ^{+0.046} _{-0.046}	-0.12 ^{+0.056} _{-0.056}	-0.06 ^{+0.04} _{-0.04}
Any	25	150	2000-2009 2015-2019	+0.34 ^{+0.044} _{-0.044}	+0.06 ^{+0.04} _{-0.04}	-0.09 ^{+0.06} _{-0.06}	-0.009 ^{+0.09} _{-0.09}	-0.11 ^{+0.047} _{-0.047}	-0.13 ^{+0.057} _{-0.057}	-0.02 ^{+0.04} _{-0.04}

Table 6: Summary of the effect of each predictor on the cognitive distance (d_a) across topic models and temporal segmentations. Values indicate the mean posterior effect size and the 95% credible interval. Significant effects are shown in bold.

Author-ship	K_0	L	Temporal segmentation	Intell. capital (diversity)	Soc. capital (diversity)	Soc. capital (power)	Stable affiliation	Academic age	Prod. (co-auth.)	Prod. (solo-auth.)
Any	20	50	2000-2009 2015-2019	+0.33 ^{+0.043} _{-0.042}	+0.11 ^{+0.04} _{-0.041}	-0.14 ^{+0.061} _{-0.061}	-0.007 ^{+0.09} _{-0.09}	-0.07 ^{+0.05} _{-0.05}	-0.1 ^{+0.06} _{-0.06}	-0.04 ^{+0.04} _{-0.04}
1st/last	20	50	2000-2009 2015-2019	+0.3 ^{+0.056} _{-0.055}	+0.1 ^{+0.053} _{-0.052}	-0.06 ^{+0.08} _{-0.08}	-0.01 ^{+0.1} _{-0.1}	-0.07 ^{+0.06} _{-0.06}	-0.12 ^{+0.076} _{-0.075}	-0.04 ^{+0.05} _{-0.05}
Any	20	50	2000-2004 2005-2009	+0.37 ^{+0.044} _{-0.044}	+0.13 ^{+0.043} _{-0.044}	-0.05 ^{+0.06} _{-0.06}	-0.04 ^{+0.1} _{-0.1}	-0.07 ^{+0.05} _{-0.05}	-0.19 ^{+0.06} _{-0.06}	-0.07 ^{+0.04} _{-0.04}
Any	20	50	2010-2014 2015-2019	+0.37 ^{+0.037} _{-0.037}	+0.09 ^{+0.04} _{-0.04}	-0.07 ^{+0.05} _{-0.05}	-0.12 ^{+0.086} _{-0.086}	-0.03 ^{+0.04} _{-0.04}	-0.21 ^{+0.05} _{-0.051}	-0.04 ^{+0.04} _{-0.04}
Any	20	50	2000-2009 2010-2019	+0.4 ^{+0.031} _{-0.031}	+0.06 ^{+0.031} _{-0.031}	-0.09 ^{+0.05} _{-0.05}	+0.01 ^{+0.07} _{-0.07}	-0.05 ^{+0.035} _{-0.035}	-0.19 ^{+0.046} _{-0.046}	-0.03 ^{+0.03} _{-0.03}
Any	20	50	2000-2009 2015-2019	+0.36 ^{+0.043} _{-0.043}	+0.11 ^{+0.04} _{-0.04}	-0.14 ^{+0.059} _{-0.06}	+0.003 ^{+0.09} _{-0.09}	-0.09 ^{+0.05} _{-0.05}	-0.14 ^{+0.056} _{-0.056}	-0.05 ^{+0.04} _{-0.04}
Any	15	50	2000-2009 2015-2019	+0.28 ^{+0.041} _{-0.041}	+0.1 ^{+0.04} _{-0.04}	-0.18 ^{+0.057} _{-0.058}	-0.04 ^{+0.09} _{-0.09}	-0.09 ^{+0.05} _{-0.05}	-0.05 ^{+0.05} _{-0.06}	-0.03 ^{+0.04} _{-0.04}
Any	25	50	2000-2009 2015-2019	+0.25 ^{+0.043} _{-0.043}	+0.12 ^{+0.041} _{-0.04}	-0.18 ^{+0.059} _{-0.058}	+0.06 ^{+0.09} _{-0.09}	-0.12 ^{+0.053} _{-0.053}	-0.06 ^{+0.06} _{-0.06}	-0.03 ^{+0.04} _{-0.04}
Any	15	150	2000-2009 2015-2019	+0.28 ^{+0.048} _{-0.047}	+0.07 ^{+0.04} _{-0.04}	-0.09 ^{+0.06} _{-0.06}	+0.03 ^{+0.09} _{-0.1}	-0.07 ^{+0.05} _{-0.05}	-0.08 ^{+0.06} _{-0.06}	-0.02 ^{+0.04} _{-0.04}
Any	20	150	2000-2009 2015-2019	+0.27 ^{+0.046} _{-0.045}	+0.04 ^{+0.04} _{-0.04}	-0.1 ^{+0.062} _{-0.061}	+0.06 ^{+0.09} _{-0.09}	-0.1 ^{+0.05} _{-0.05}	-0.09 ^{+0.06} _{-0.06}	-0.04 ^{+0.04} _{-0.04}
Any	25	150	2000-2009 2015-2019	+0.26 ^{+0.047} _{-0.047}	+0.07 ^{+0.043} _{-0.043}	-0.15 ^{+0.063} _{-0.064}	+0.02 ^{+0.09} _{-0.09}	-0.08 ^{+0.05} _{-0.05}	-0.08 ^{+0.06} _{-0.06}	-0.003 ^{+0.04} _{-0.04}

Table 7: Summary of the effect of each predictor on the probability of having entered a research area across topic models and temporal segmentations. Values indicate the mean posterior effect size and the 95% credible interval. Significant effects are shown in bold.

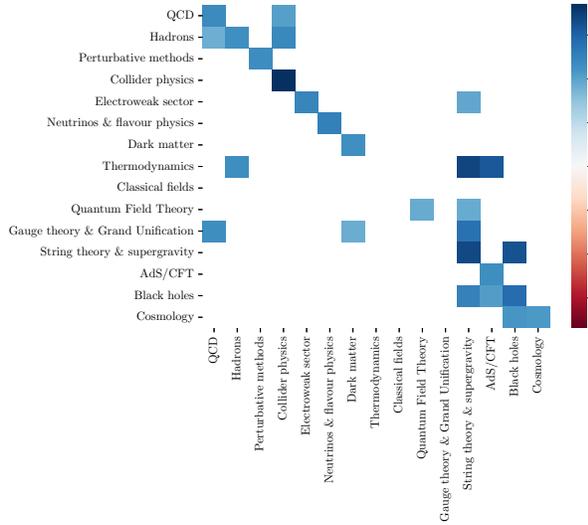
Author-ship	K_0	L	Temporal segmentation	Intell. capital (diversity)	Soc. capital (diversity)	Soc. capital (power)	Stable affiliation	Academic age	Prod. (co-auth.)	Prod. (solo-auth.)
Any	20	50	2000-2009 2015-2019	+0.2 ^{+0.11} _{-0.11}	+0.22 ^{+0.1} _{-0.1}	+0.006 ^{+0.1} _{-0.1}	-0.19 ^{+0.22} _{-0.22}	+0.04 ^{+0.12} _{-0.11}	-0.07 ^{+0.1} _{-0.1}	-0.05 ^{+0.1} _{-0.09}
1st/last	20	50	2000-2009 2015-2019	+0.25 ^{+0.14} _{-0.13}	+0.19 ^{+0.13} _{-0.13}	+0.21 ^{+0.19} _{-0.19}	+0.12 ^{+0.29} _{-0.29}	-0.07 ^{+0.1} _{-0.1}	-0.33 ^{+0.18} _{-0.18}	-0.0006 ^{+0.1} _{-0.1}
Any	20	50	2000-2004 2005-2009	+0.16 ^{+0.11} _{-0.11}	+0.2 ^{+0.11} _{-0.11}	+0.01 ^{+0.15} _{-0.15}	-0.007 ^{+0.2} _{-0.2}	-0.06 ^{+0.1} _{-0.1}	-0.14 ^{+0.14} _{-0.14}	-0.05 ^{+0.1} _{-0.1}
Any	20	50	2010-2014 2015-2019	+0.24 ^{+0.095} _{-0.095}	+0.18 ^{+0.091} _{-0.091}	-0.06 ^{+0.1} _{-0.1}	+0.05 ^{+0.2} _{-0.21}	+0.07 ^{+0.1} _{-0.1}	-0.18 ^{+0.12} _{-0.12}	-0.02 ^{+0.09} _{-0.09}
Any	20	50	2000-2009 2010-2019	+0.27 ^{+0.083} _{-0.083}	+0.27 ^{+0.076} _{-0.076}	-0.01 ^{+0.1} _{-0.1}	-0.08 ^{+0.2} _{-0.2}	+0.04 ^{+0.09} _{-0.09}	-0.16 ^{+0.11} _{-0.11}	-0.004 ^{+0.07} _{-0.07}
Any	20	50	2000-2009 2015-2019	+0.24 ^{+0.11} _{-0.11}	+0.2 ^{+0.1} _{-0.1}	+0.03 ^{+0.15} _{-0.14}	-0.07 ^{+0.2} _{-0.2}	-0.01 ^{+0.1} _{-0.1}	-0.19 ^{+0.13} _{-0.14}	-0.06 ^{+0.1} _{-0.09}
Any	15	50	2000-2009 2015-2019	+0.13 ^{+0.11} _{-0.1}	+0.13 ^{+0.096} _{-0.095}	-0.005 ^{+0.1} _{-0.1}	-0.03 ^{+0.2} _{-0.2}	-0.04 ^{+0.1} _{-0.1}	-0.16 ^{+0.13} _{-0.13}	+0.002 ^{+0.09} _{-0.09}
Any	25	50	2000-2009 2015-2019	+0.18 ^{+0.12} _{-0.12}	+0.17 ^{+0.11} _{-0.11}	-0.009 ^{+0.2} _{-0.2}	-0.07 ^{+0.2} _{-0.2}	-0.04 ^{+0.1} _{-0.1}	-0.15 ^{+0.15} _{-0.15}	-0.02 ^{+0.1} _{-0.1}
Any	15	150	2000-2009 2015-2019	+0.12 ^{+0.1} _{-0.1}	+0.19 ^{+0.095} _{-0.095}	-0.08 ^{+0.1} _{-0.1}	-0.02 ^{+0.2} _{-0.2}	-0.02 ^{+0.1} _{-0.1}	-0.04 ^{+0.1} _{-0.1}	-0.06 ^{+0.09} _{-0.09}
Any	20	150	2000-2009 2015-2019	+0.23 ^{+0.11} _{-0.11}	+0.07 ^{+0.1} _{-0.1}	+0.12 ^{+0.15} _{-0.15}	-0.06 ^{+0.2} _{-0.2}	-0.08 ^{+0.1} _{-0.1}	-0.21 ^{+0.14} _{-0.14}	-0.05 ^{+0.095} _{-0.093}
Any	25	150	2000-2009 2015-2019	+0.31 ^{+0.13} _{-0.13}	+0.01 ^{+0.1} _{-0.1}	-0.09 ^{+0.2} _{-0.2}	-0.14 ^{+0.25} _{-0.25}	+0.06 ^{+0.1} _{-0.1}	+0.001 ^{+0.2} _{-0.2}	-0.02 ^{+0.1} _{-0.1}

Table 8: Summary of the effect of each predictor on the probability of having exited a research area across topic models and temporal segmentations. Values indicate the mean posterior effect size and the 95% credible interval. Significant effects are shown in bold.

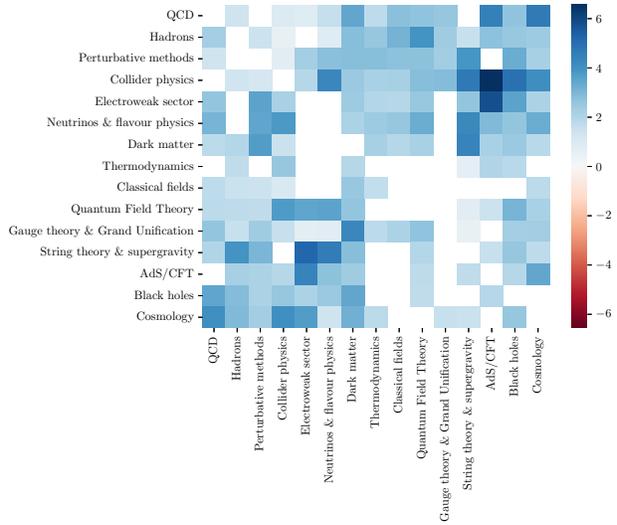
Author-ship	K_0	L	Temporal segmentation	Intell. capital (diversity)	Soc. capital (diversity)	Soc. capital (power)	Stable affiliation	Academic age	Prod. (co-auth.)	Prod. (solo-auth.)
Any	20	50	2000-2009 2015-2019	+1 ^{+0.14} _{-0.14}	+0.04 ^{+0.1} _{-0.1}	-0.03 ^{+0.2} _{-0.2}	+0.04 ^{+0.2} _{-0.2}	-0.21 ^{+0.12} _{-0.12}	-0.28 ^{+0.15} _{-0.14}	-0.02 ^{+0.1} _{-0.1}
1st/last	20	50	2000-2009 2015-2019	+1.1 ^{+0.18} _{-0.18}	+0.01 ^{+0.1} _{-0.1}	-0.12 ^{+0.2} _{-0.2}	-0.05 ^{+0.3} _{-0.3}	-0.15 ^{+0.16} _{-0.16}	-0.12 ^{+0.19} _{-0.19}	-0.01 ^{+0.1} _{-0.1}
Any	20	50	2000-2004 2005-2009	+1 ^{+0.15} _{-0.14}	+0.04 ^{+0.1} _{-0.1}	-0.18 ^{+0.16} _{-0.16}	-0.07 ^{+0.3} _{-0.3}	-0.18 ^{+0.13} _{-0.13}	-0.25 ^{+0.15} _{-0.15}	+0.04 ^{+0.1} _{-0.1}
Any	20	50	2010-2014 2015-2019	+0.89 ^{+0.12} _{-0.11}	+0.04 ^{+0.09} _{-0.1}	+0.1 ^{+0.1} _{-0.1}	-0.22 ^{+0.22} _{-0.22}	-0.04 ^{+0.1} _{-0.1}	-0.34 ^{+0.13} _{-0.13}	-0.09 ^{+0.09} _{-0.09}
Any	20	50	2000-2009 2010-2019	+0.88 ^{+0.093} _{-0.09}	+0.16 ^{+0.077} _{-0.076}	-0.01 ^{+0.1} _{-0.1}	+0.07 ^{+0.2} _{-0.2}	-0.06 ^{+0.09} _{-0.08}	-0.32 ^{+0.11} _{-0.11}	-0.09 ^{+0.07} _{-0.07}
Any	20	50	2000-2009 2015-2019	+0.91 ^{+0.13} _{-0.13}	-0.05 ^{+0.1} _{-0.1}	-0.07 ^{+0.2} _{-0.2}	+0.001 ^{+0.2} _{-0.2}	-0.16 ^{+0.12} _{-0.12}	-0.28 ^{+0.14} _{-0.14}	-0.002 ^{+0.1} _{-0.1}
Any	15	50	2000-2009 2015-2019	+1.1 ^{+0.14} _{-0.14}	-0.003 ^{+0.1} _{-0.1}	-0.15 ^{+0.15} _{-0.15}	-0.19 ^{+0.23} _{-0.23}	-0.05 ^{+0.1} _{-0.1}	-0.21 ^{+0.14} _{-0.14}	-0.07 ^{+0.1} _{-0.1}
Any	25	50	2000-2009 2015-2019	+1.1 ^{+0.16} _{-0.16}	-0.07 ^{+0.1} _{-0.1}	+0.17 ^{+0.17} _{-0.17}	-0.06 ^{+0.3} _{-0.3}	-0.21 ^{+0.14} _{-0.14}	-0.39 ^{+0.15} _{-0.16}	+0.09 ^{+0.1} _{-0.1}
Any	15	150	2000-2009 2015-2019	+1.3 ^{+0.15} _{-0.14}	+0.05 ^{+0.1} _{-0.1}	-0.11 ^{+0.15} _{-0.15}	+0.04 ^{+0.2} _{-0.2}	-0.06 ^{+0.1} _{-0.1}	-0.23 ^{+0.15} _{-0.15}	+0.04 ^{+0.1} _{-0.1}
Any	20	150	2000-2009 2015-2019	+1.1 ^{+0.14} _{-0.14}	+0.09 ^{+0.1} _{-0.1}	-0.09 ^{+0.2} _{-0.2}	+0.09 ^{+0.2} _{-0.2}	-0.11 ^{+0.12} _{-0.12}	-0.21 ^{+0.15} _{-0.15}	-0.06 ^{+0.1} _{-0.1}
Any	25	150	2000-2009 2015-2019	+1.1 ^{+0.17} _{-0.17}	+0.1 ^{+0.1} _{-0.1}	+0.16 ^{+0.18} _{-0.18}	+0.24 ^{+0.27} _{-0.27}	-0.14 ^{+0.14} _{-0.14}	-0.44 ^{+0.16} _{-0.16}	-0.05 ^{+0.1} _{-0.1}

A.9 Trajectory model parameters evaluated on a different time period

To further assess the robustness of the findings, the effect of intellectual and social capital on individual trajectories is measured on a different temporal segmentation (2000-2004 to 2005-2009). We make similar findings: the concentration of intellectual capital in one area promotes either commitment to this research area (or transfers in related areas). Social capital, on the other hand, matters increasingly as cognitive distance increases.



(a) Effect of intellectual capital ($\gamma_{kk'}$).



(b) Effect of social capital ($\delta_{kk'}$).

Figure 19: **Effect of intellectual and social capital on transfers across research areas.** Rows represent research areas of origin and columns represent target areas. Effect sizes are expressed in log-odds ($\log [\theta_{akk'}/(1 - \theta_{akk'})]$, where $\theta_{akk'}$ is the fraction of attention to k redirected to k') per unit of intellectual or social capital in the target research area k' . Effects that are not significant (at the 95% credible level) are displayed in white for purposes of clarity.

A.10 Replication corpus

For purposes of testing and replication, certain analyses have been reproduced on the ACL anthology corpus of Computational Linguistics research .

The transfers of attention are shown in Figure 20. Compared to the high-energy physics corpus, it features significant disruptions (e.g., the emergence of new topics, such as “deep learning”, “sentiment analysis” and “embeddings & pre-trained models”).

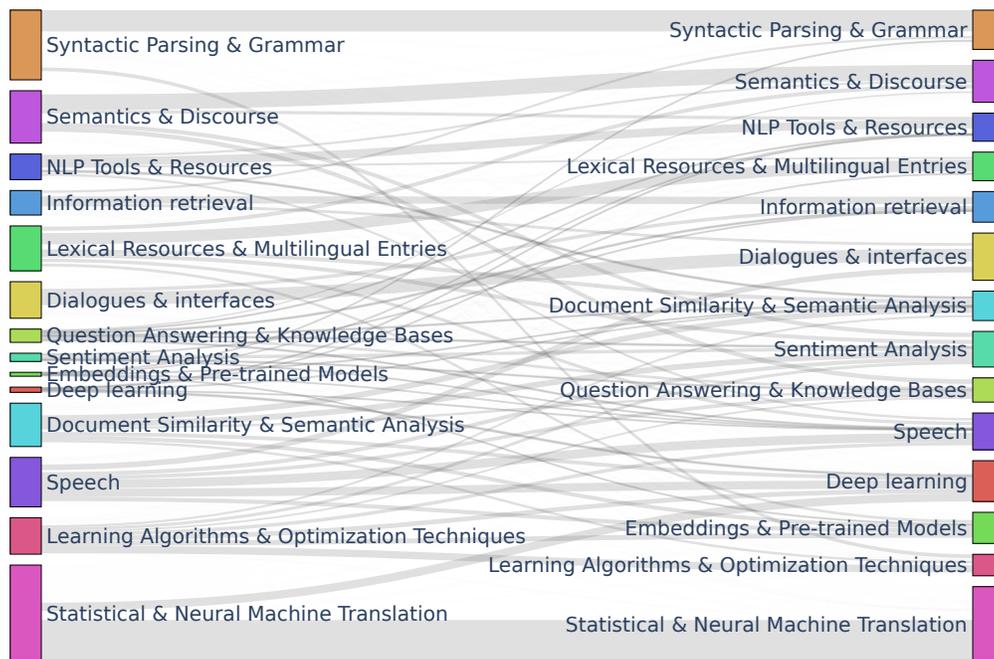


Figure 20: **Aggregate transfers of attention across research areas in the ACL anthology corpus of Computational Linguistics, between 2002-2011 (to the left) and 2012-2022 (to the right).** Widths of flows are proportional to $\sum_a X_{ak} \theta_{akk'}$. Transfers less frequent than expected by chance alone are transparent.