NBER WORKING PAPER SERIES

SPECIALIZATION TRENDS IN ECONOMICS RESEARCH: A LARGE-SCALE STUDY USING NATURAL LANGUAGE PROCESSING AND CITATION ANALYSIS

Sebastian Galiani
Ramiro H. Gálvez
Ian Nachman

Working Paper 31295
http://www.nber.org/papers/w31295

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2023, revised September 2024

Specialization Trends in Economics Research: A Large-Scale Study Using Natural Language Processing and Citation Analysis
Sebastian Galiani, Ramiro H. Gálvez, and Ian Nachman
NBER Working Paper No. 31295
June 2023, revised September 2024
JEL No. A1

ABSTRACT

This article conducts a comprehensive analysis of specialization trends within and across fields of economics research. We collect data on 24,273 articles published between 1970 and 2016 in general research economics outlets and employ machine learning techniques to enrich the collected data. Results indicate that theory and econometric methods papers are becoming increasingly specialized, with a narrowing scope and steady or declining citations from outside economics and from other fields of economics research. Conversely, applied papers are covering a broader range of topics, receiving more extramural citations from fields like medicine, and psychology. Trends in applied theory articles are unclear.

Sebastian Galiani
Department of Economics
University of Maryland
3105 Tydings Hall
College Park, MD 20742
and NBER
sgaliani@umd.edu

Ramiro H. Gálvez
Universidad Torcuato Di Tella
Av. Figueroa Alcorta 7350, (C1428BIJ)
Buenos Aires
Argentina
rgalvez@utdt.edu

Ian Nachman
ian_nachman@brown.edu

Specialization Trends in Economics Research: A Large-Scale Study Using Natural Language Processing and Citation Analysis

Sebastian Galiani† — *University of Maryland & NBER*

Ramiro H. Gálvez‡ — *Universidad Torcuato Di Tella*

Ian Nachman§ — *Brown University*

August 26, 2024

Abstract

This article conducts a comprehensive analysis of specialization trends within and across fields of economics research. We collect data on 24,273 articles published between 1970 and 2016 in general research economics outlets and employ machine learning techniques to enrich the collected data. Results indicate that theory and econometric methods papers are becoming increasingly specialized, with a narrowing scope and steady or declining citations from outside economics and from other fields of economics research. Conversely, applied papers are covering a broader range of topics, receiving more extramural citations from fields like medicine, and psychology. Trends in applied theory articles are unclear. (JEL A11, A14)

Keywords: *Fields of economics research*; *Specialization trends*; *Machine learning*; Natural language processing; *Citation analysis*

1. Introduction

Knowledge is commonly divided into fields of study or academic disciplines (e.g., economics, psychology, mathematics, etc., see Krishnan, 2009). Fields are typically defined and circumscribed by the journals that publish their research, their learned societies, and the academic departments to which their practitioners belong. Usually, a field of study focuses on a series of research topics. However, defining a field of study solely by the research topics it addresses is inadequate, as, in many cases, the same research topic is addressed by multiple fields of study by means of different research strategies/tools. A well-documented fact is that, as time passes by and knowledge accumulates, fields of study tend to become more specialized, to address increasingly complex research questions and to develop new, tailored research tools (see Ramaley, 1930; Ziman, 1987; Wray, 2005; *inter alia*). Although specialization is believed to lead to deeper and more precise knowledge, it has also been identified as a factor creating barriers to communication and collaboration across different fields of study (see, for example, Becker and Murphy, 1992; Walsh and Maloney, 2007; Anderson and Richards-Shubik, 2022).[1]

Within a field of study, knowledge is also typically divided into subfields, each often having its own dedicated journals and learned societies.[2] Researchers within a subfield tend to focus their research on a subset of the topics addressed by the parent field. However, another common division within fields involves categorizing them into a theoretical branch and an experimental/applied branch. For instance, physics is commonly segmented into theoretical physics and experimental physics (Duhem, 1976). Theoretical physicists employ advanced mathematical models to formulate theories

† sgaliani@umd.edu; ‡ rgalvez@utdt.edu; § ian_nachman@brown.edu.

[1] William Dunham, in his book "*The mathematical universe*" (Dunham, 1994), presents the reader with an exercise that illustrates this trend. He notes that any mathematician asked to name history's three or four most influential mathematicians would surely include Sir Isaac Newton in their list. Similarly, any physicist asked to name history's three or four most influential physicists would also include Newton. Dunham concludes that, although this is extraordinarily rare, it occurs in part because Newton worked at a time before "*unscalable walls were erected between disciplines*." He also notes that, at the time, subjects such as optics, astronomy, and mechanics were in fact treated as branches of mathematics.

[2] For example, medicine is divided into anatomy, physiology, pathology, etc.; chemistry is divided into analytical chemistry, organic chemistry, inorganic chemistry, etc.; economics is divided into labor economics, behavioral economics, economic development; etc.

explaining physical systems, while experimental physicists design and conduct experiments to test these theories. Similar divisions exist in other fields, such as theoretical psychology and experimental psychology, mathematical biology and experimental biology, theoretical linguistics and experimental linguistics, and so forth. Ideally, there should be a feedback loop between the two branches, with the theoretical branch proposing new theories for testing, and the empirical branch providing novel data essential for developing, refuting, or validating these theories (see Putnam, 1979; Reeves et al., 2008).

Economics is no exception, and a clear division between theoretical and applied economics exists (see Backhouse and Cherrier, 2014; Hamermesh, 2018; Angrist et al. 2017; *inter alia*). Specialized journals exist for both theory and applied economics, applied/theory oriented postgraduate programs exist, and, some universities even have dedicated applied economics departments. What is less clear is how these branches interact with each other and with other fields of study. Do they provide feedback to each other? Or, as is often believed to be the case across disciplines, are they becoming so specialized that collaboration and interaction are being precluded? Focusing on what we refer to as "fields of economics research" (*applied*, *applied theory*, *econometric methods*, and *theory*), we empirically address this issue by analyzing long-term trends in the contents of, and the citations received and generated by, a large and representative corpus of economics research articles.[3] Concretely, we look for patterns that indicate specialization within a field of economics research, such as a narrowing of the topics it covers in a way that other fields do not, and patterns suggesting a decrease in citations from outside the field.[4]

The research strategy we implement in this article is as follows. We first collect a large sample of economics research articles published in top outlets. The sample covers a wide and representative set of research topics/strategies and spans the period 1970-2016. Then, for each article in our sample we collect extensive detailed data (e.g., metadata, word counts, data on the articles citing them, data on the articles they cite, and so on). Having collected these data, we enhance/enrich it using modern machine learning (ML) and natural language processing (NLP) techniques. For example, 1) based on the articles' word counts and on a set hand-labelled field of economics research tags obtained from previous studies, we predict the field of economics research to which each article in our sample belongs; 2) based solely on the articles' word counts, we identify the topics addressed by each article in our dataset; and 3) based on the titles and abstracts of articles citing or being cited by our sample, we predict the field of economics research to which these citations or references belong. The resulting data enable us to examine previously undocumented trends characterizing specialization in economics.[5] This includes, for example, assessing the extent to which applied papers cite or reference theory articles, the degree to which different fields of economics research have concentrated their research efforts on a few topics, which topics have gained more presence in theory or applied papers in the last decades, and more.

Our analysis reveals significant disparities in the content, citations, and references of articles across fields of economics research.[6] Certain fields of economics research have demonstrated a

[3] As we will detail in Section 3.3, we define a field of economics research based on the methodological techniques/strategies used by its articles to address research topics/questions. As stated in Hamermesh (2013), "*subject does not imply method*," as a given economics research topic (e.g., corruption) is commonly addressed by means of different research strategies (e.g., by running a lab experiment or by means of a theoretical model). In fact, it could easily be argued that researchers require less training to migrate from one research topic to another while using the same research methodology than to study the same topic by means of different research methodologies.

[4] Specialization may manifest along dimensions beyond the content of articles and their received citations. An intriguing avenue for research involves exploring changes in authorship over time. It is plausible that authors contributing to articles across various fields of economics research may be less common today than in the past, suggesting a trend toward specialization. However, a comprehensive analysis of such behaviors poses several challenges, including accurately identifying authors, addressing co-authored papers, and determining an adequate sample of articles for such an analysis. Given these considerations and for the sake of parsimony, this work does not delve into the study of articles' authorship.

[5] The data and scripts that support the findings of this study are openly available in openICPSR at https://doi.org/10.3886/E198921, reference number 198921 (Galiani, Gálvez, and Nachman 2024).

[6] An alternative approach could have involved exploring variations in specialization trends across JEL codes (a standardized classification system by the Journal of Economic Literature for organizing academic research articles), which divide economics

growing trend towards specialization, while others exhibited contrasting patterns. Specifically, theory and econometric methods have shown a narrowing focus on specific research topics since the 1990s, indicating a tendency towards specialization. Theory articles have experienced a significant increase in topics related to formal mathematical proofs and game theory, while econometric methods articles have shown a pronounced rise in topics related to computational statistics, estimators' asymptotic properties, and estimators' bounds. In contrast to applied papers, these fields do not exhibit rising trends in extramural citations (i.e., citations from other disciplines) and in citations from other fields within economics research (in the case of econometric methods, it shows a declining trend). These patterns also indicate a higher degree of specialization.

Applied papers have expanded their coverage to include diverse topics. Applied articles have seen a pronounced rise in topics related to impact analysis, causal analysis, and experimental economics. Over time, applied articles began to receive a higher proportion of citations from external fields, especially from disciplines such as medicine, psychology, law, and to a certain extent, education. By 2016, applied ranked among the most or second most cited field by any other field of economics research. Overall, these patterns indicate that applied papers are becoming more multidisciplinary. The case of applied theory articles is less conclusive. While they cover a broader range of topics (similar to applied papers), there has been no significant increase in extramural citations or citations from other fields of economics research (as observed with theory articles). In fact, applied theory articles receive the smallest share of extramural citations. These contrasting patterns pose a challenge in determining their specialization status.

The rest of the article is structured as follows. Section 2 provides a review of the previous literature. Section 3 offers a detailed description of the data sources utilized in this study. In Section 4, we explain the filtering, processing and enrichment of the collected data through the application of ML/NLP techniques. Section 5 presents the main findings of our analysis. Finally, Section 6 concludes by summarizing our results and discussing practical implications.

2. Literature review

The division of an academic discipline into empirical and theoretical branches is a widely recognized and extensively studied phenomenon across multiple disciplines (see, Poincaré, 1902; Dirac, 1940; Koch, 1973; Harrow, 1979; Longo and Soto, 2016; Phillips et al., 2021; *inter alia*). Previous papers within the field of economics have also examined differences along this dimension. Hamermesh (2013) provides data on articles published in three general research economics journals for one year in each decade from the 1960s to the 2010s. The author documents that these journals are publishing fewer papers that represent pure theory, fewer papers based on publicly available datasets, and more empirical studies relying on data collected by their respective authors or on laboratory and field experiments. In Hamermesh (2018) a field of economics research tag is assigned to 439 articles published in 1974-75 and to 497 articles published in 2007-08. The author documents that empirical and experimental articles receive more citations than economic theory or econometric theory articles do. Chiappori and Levitt (2003) explore whether theoretical economic research succeeds in influencing the path of empirical microeconomic research. They examine a dataset consisting of empirical microeconomic papers published between 1999 and 2001. They find that the set of theoretical papers cited as a primary motivation for empirical research projects is surprisingly diverse and that very few theoretical papers have much of an influence on applied microeconomic papers. Biddle and Hamermesh (2017) code the content of all applied microeconomic articles published in the top economics journals in the years 1951-1955, 1974-1975, and 2007-2008. They identify three separate

into cohesive fields in relation to the topics covered. However, as previously mentioned, our primary focus is on understanding distinctions among categories determined by the strategies and skills employed in writing economics articles (it is worth noting that papers sharing the same JEL code might employ entirely different research strategies). Future research endeavors could investigate the prevalence of the trends we have identified across articles categorized with different JEL codes.

stages in the post-World War II history of applied microeconomic research: a nonmathematical period, a period of consensus, and a period referred to as the "experimentalist paradigm." They also report a partial abandonment of theory by applied microeconomists. Backhouse and Cherrier (2014) examine the historical winners of the John Bates Clark medal. They conclude that economics has become increasingly applied over time, with applied work now being accorded higher status compared to pure theory.

Anauati, Galiani, and Gálvez (2016) characterize how the life cycles in yearly citations differ across four fields of economic research (applied, applied theory, econometric methods and theory — the same categories used in this article). They assign a field of economics research tag to 9,672 economics research articles published in top journals between 1970 and 2000 and collect detailed citation data for them. They document strong differences in citation patterns across fields of economics research. For example, applied and applied theory articles receive higher numbers of citations per year and have the longer citation lifespans. Using the same categories of economics research fields, Anauati, Galiani, and Gálvez (2020) study how citation patterns differ among journal tiers in economics (Top 5, non-Top 5 general research, and Top Field). They assign fields of economics research tags to 6,083 articles published across different journal tiers. Among their findings, they report that differences in citation performance across journal tiers are significantly influenced by the fields of economics research covered in the articles. In this article, we use data from Anauati, Galiani, and Gálvez (2016) and from Anauati, Galiani, and Gálvez (2020) to train ML/NLP models that assign a field of economics research tag to a much larger sample of economics research articles (see Section 3.3 and Section 4).

The two papers most closely related to ours are Angrist et al. (2017) and Angrist et al. (2020). In both articles, the authors utilize ML/NLP techniques to assign specific field of economics research tags to a large sample of economic research articles. Both articles examine three distinct fields of economics research: *empirical*, *theory*, and *econometrics* (although they use the term "style" instead of fields, as we do). In Angrist et al. (2017), the authors examine the purported shift in economics from theoretical to empirical research. They also document that empirical work is now cited more often than theoretical work. In Angrist et al. (2020), the authors look at the influence of economic research through the lens of *extramural citations* (i.e., citations received from fields of study other than economics). The authors document a clear rise in the extramural influence of economics research, along with a growing tendency for economics to reference and draw insights from other social sciences. They observe that a significant portion of the increase in economics' external influence can be attributed to the rise in citations to empirical work. Additionally, they document an increase in the proportion of citations attributed to empirical papers within the field of economics.

The present paper differs from previous studies in two important ways. First, it primarily focuses on documenting specialization trends within fields of economics research. While effectively recording patterns that have been previously studied (such as the shares of extramural citations received by different fields of economics research), we go beyond existing literature by uncovering previously unexplored trends. Second, our paper introduces several methodological innovations. On one hand, we analyze new data sources that provide a higher level of granularity compared to those examined in previous studies. This enables us to study a larger, richer, and more representative corpus, spanning a longer period than most previous literature.[7] On the other hand, the collected data allows us to utilize ML/NLP techniques that have not been employed in previous literature studying corpora of economics research articles. For instance, to the best of our knowledge, this study represents the first application of topic analysis on such a vast and comprehensive corpus of economics research articles.

[7] Three papers studying article content for a period and sample of journals as large as the one we analyze are Kosnik (2015), Kosnik (2016), and Hamermesh and Kosnik (2023). Angrist et al. (2017) and Angrist et al. (2020) study the period from 1980 to 2015, which is slightly shorter than the period we analyze (1970-2016). The remaining literature either studies shorter periods or investigates similarly sized periods, but with significant sampling from the population of economics articles published in the journals they analyze.

3. Data sources

We utilize data from three sources: 1) *Constellate*, 2) *Semantic Scholar*, and 3) a set of fields of economics research tags taken from Anauati, Galiani, and Gálvez (2016) and Anauati, Galiani, and Gálvez (2020). Below, we provide further details on each of these sources.

3.1. Constellate

Launched in 2021, Constellate is a text analytics service provided by ITHAKA (the not-for-profit organization managing both JSTOR and Portico).[8] Constellate includes an online dataset builder which allows to retrieve data from articles satisfying a series of criteria (e.g., articles that have been published in a given journal, that have been published in a given year, etc.). The downloaded data includes metadata for each article that meets the criteria, such as title, digital object identifier (DOI), author names, publication date, journal of publication, and more. Importantly, it also includes the counts of unigrams, bigrams, and trigrams associated with each article retrieved.[9]

Constellate covers only a subset of all published economics research journals. For example, data for articles from the journal *Economic Theory* cannot be downloaded from Constellate. Taking this into account, we opted to fetch Constellate data for as many general research journals as were studied in Anauati, Galiani, and Gálvez (2016) and Anauati, Galiani, and Gálvez (2020). That is, we downloaded data for all articles published from 1970 up to and including 2016 in the so-called economics *Top 5* journals (*The American Economic Review*, *Journal of Political Economy*, *The Quarterly Journal of Economics*, *Econometrica*, and *The Review of Economic Studies*) and for all articles published from 1970 up to and including 2016 in a set of *non-Top 5* general research journals (*The Economic Journal*, *International Economic Review*, *Economic Inquiry*, and *The Review of Economics and Statistics*).[10] Taken together, these journals extensively cover all fields of economics research and have been the subject of study in numerous articles similar in scope to ours (e.g., Heckman and Moktan, 2020; Angrist et al., 2020; Card and DellaVigna, 2013; Hamermesh, 2018, etc).[11]

Solely for training the NLP models described in Section 4.1 and Section 4.2, we downloaded Constellate data for articles published from 1970 to 2016 in the *Journal of Law & Economics*, the *Journal of Labor Economics* and the *RAND Journal of Economics*. These are the only "top field" journals considered in Anauati, Galiani, and Gálvez (2020) that are available for download from Constellate throughout most of the entire period.

After downloading this data, we applied a series of filters. First, we retained documents that had the tag "research-article" or "article" in the "docSubType" field, removing those classified as miscellaneous by Constellate. Second, we dropped all documents for which the field “creator” (which lists the articles’ authors) was empty or missing. Third, in cases where two articles shared titles, and one was retrieved from JSTOR while the other from Portico, we retained only the JSTOR version. Fourth, through an analysis of their titles, we excluded specific articles identified as not being proper research articles (e.g., replies, rejoinders, errata, editor's reports, software reviews, etc.). Finally, we excluded documents with the field “title” missing. After applying these filters, the Constellate data consisted of 34,623 documents.

[8] https://constellate.org/.

[9] Unigram count stands for the time a given word/token appears in a document. Bigram counts stands for the times a given sequence of two words/tokens (e.g., "your homework") appears in a document. Trigram counts is the equivalent, but for sequences of three words (e.g., "turn your homework"). See Jurafsky and Martin (2009).

[10] We downloaded data up to and including 2016, as Constellate did not include data on newer articles at the time of download.

[11] We recognize that by concentrating on general interest articles, there is a risk that our observations may not necessarily reflect what economists are actively engaged in, but rather the view of a narrow group of journal editors (May et al. 2021). Conclusions derived from this article should take this fact into consideration. Future research endeavors could delve into examining the prevalence of the trends we report across a more extensive array of journals, encompassing a broader set of topics and authors.

3.2. *Semantic Scholar*

As we will describe in detail in Section 4.2, we downloaded Semantic Scholar data for all articles included in the Constellate data. Semantic Scholar is an artificial intelligence powered research tool for scientific literature developed by the Allen Institute for AI (Kinney et al., 2023).[12] Launched in 2015, the tool provides an academic search engine (accessible through an API), and, by early 2023, indexed more than 207 million published scientific research articles, books, and preprints (see Matthews 2021). The Semantic Scholar's corpus includes documents from all fields of research and is considered comparable in coverage to that of Scopus and the Web of Science (see Martín-Martín et al., 2021).[13]

After conducting a search, Semantic Scholar provides various data. This includes metadata for all retrieved articles (title, publication date, identifiers, authors, fields of study tags, etc.), abstracts of the retrieved articles, and comprehensive information on the articles cited by the retrieved articles (i.e., on their references). Additionally, Semantic Scholar provides detailed information on all articles that cite the retrieved articles, including their publication year, abstract, and fields of study tags.

The fields of study tags are labels that categorize articles into one or more of the following disciplines: economics, history, philosophy, computer science, business, medicine, physics, political science, mathematics, psychology, sociology, geology, environmental science, law, biology, engineering, education, art, geography, agricultural and food sciences, materials science, chemistry, and linguistics.[14] Importantly, in this study, we consider an article to be an economics article if "economics" is listed among its Semantic Scholar fields of study tags, regardless of whether any other field is also listed.

3.3. *Fields of economics research hand-labelled tags*

In Anauati, Galiani, and Gálvez (2016), a total of 9,672 articles published in the Top 5 journals from 1970 to 2000 were manually tagged as belonging to one of the following four fields of economics research: 1) applied, 2) applied theory, 3) econometric methods, and 4) theory. The field of economics research categories used in that article aimed to capture the research strategy employed and the skills required when writing an economics article. The criteria used to assign a paper to a field were as follows: 1) *Applied* articles have an empirical or applied motivation. They rely on the use of econometric or statistical methods as a basis for analyzing empirical data, although they may deal with simple models that provide a theoretical framework for the analysis. This category also includes papers that do not use sophisticated econometric methods but instead use descriptive statistics to analyze, for example, given features of an economy, and in which the empirical section figures as the central element. 2) *Applied theory* articles construct a theoretical model to explain a fact, with the empirical analysis serving as a supplementary aspect rather than the primary focus of the paper. These articles typically have limited utilization of econometric or statistical analyses, though they may employ simulations (even with empirical data) or refine other techniques to test the implications of the models.[15] 3) *Econometric methods* articles develop econometric or statistical methodologies. They also include papers that develop methodologies for collecting data and that address issues of identification, data aggregation, or optimization techniques. 4) *Theory* articles lack an empirical section; typically, they approach a topic through modeling and extensive use of formal mathematics and logic. These articles may incorporate a numerical example or a simple model calibration using theoretical data to illustrate the proposed

[12] https://allenai.org/.
[13] This comes mainly from the fact that Semantic Scholar incorporated data from Microsoft Academic available via the Microsoft Academic Graph (see Martín-Martín et al., 2021).
[14] For details on how these tags are assigned, see https://blog.allenai.org/9d2f641949e5.
[15] Note that according to the criteria followed in Angrist et al. (2020), these articles would be tagged as "empirical" articles, as the authors claim that their definition of "empirical" should be understood as "*not purely theoretical*."

model or to analyze its comparative statics. For additional details on the classification criteria, refer to Anauati, Galiani, and Gálvez (2016).

In addition to the tags assigned to the articles mentioned earlier, we also utilize the field of economics research tags assigned to 549 articles published in the Top 5 journals in 2005 and 2010. These tags were initially assigned manually during the early stages of the tagging process conducted in Anauati, Galiani, and Gálvez (2016), but were ultimately not used in their final analysis.

In Anauati, Galiani, and Gálvez (2020), the same criteria used to assign articles to fields of economics research were applied to tag a series of additional articles published from 1992 to 1996 in non-Top 5 general research journals and top field journals. We reuse these tags for all articles published in the journals analyzed in Anauati, Galiani, and Gálvez (2020) that are also available in Constellate throughout most of the entire period. This adds a total of 1,681 article tags to our dataset.

After excluding seven articles identified as irrelevant or duplicates, we collected a total of 11,895 economics research field tags. We will refer to these hand-labeled tags as the *training tags* and to the collection of these articles as the *manually tagged dataset*. Table 1 displays the distribution of training tags across journals in the manually tagged dataset.

Table 1. Training tags distribution across journals in the manually tagged dataset

	Applied	Applied theory	Econometric methods	Theory
American Economic Review	990 (36.65%)	359 (13.29%)	47 (1.74%)	1305 (48.32%)
Econometrica	172 (7.14%)	155 (6.44%)	786 (32.64%)	1295 (53.78%)
Economic Inquiry	141 (64.38%)	41 (18.72%)	0 (0%)	37 (16.89%)
Economic Journal	181 (50.42%)	46 (12.81%)	23 (6.41%)	109 (30.36%)
International Economic Review	19 (7.28%)	34 (13.03%)	22 (8.43%)	186 (71.26%)
Journal of Labor Economics	94 (70.15%)	17 (12.69%)	1 (0.75%)	22 (16.42%)
Journal of Law & Economics	85 (85.86%)	5 (5.05%)	0 (0%)	9 (9.09%)
Journal of Political Economy	844 (40.48%)	286 (13.72%)	14 (0.67%)	941 (45.13%)
Quarterly Journal of Economics	516 (34.42%)	148 (9.87%)	11 (0.73%)	824 (54.97%)
RAND Journal of Economics	75 (38.27%)	11 (5.61%)	2 (1.02%)	108 (55.1%)
Review of Economic Studies	128 (8.41%)	129 (8.48%)	140 (9.2%)	1125 (73.92%)
Review of Economics and Statistics	320 (77.67%)	21 (5.1%)	66 (16.02%)	5 (1.21%)
Total	3565 (29.97%)	1252 (10.53%)	1112 (9.35%)	5966 (50.16%)

Note. Relative frequency by row in parentheses.

4. Data construction

This section serves two purposes. First, it offers a detailed walkthrough of the processing, enrichment, and merging of the data presented in Section 3 to construct our final datasets. Second, it provides a brief introduction to the NLP techniques employed in this article, encompassing both traditional and state-of-the-art methods.

4.1. Assigning a field of economics research tag to all of the Constellate articles

After downloading the Constellate detailed data (Section 3.1) and collecting hand-labelled fields of economics research tags (Section 3.3), the first step in building our final datasets is to assign each Constellate article a field of economics research tag. Since articles with manually-assigned tags represent a subset of the Constellate data, the process of assigning tags to all Constellate articles is more complex than a straightforward merge of the two datasets. In this section, we provide a detailed explanation of the approach we employed to predict field of economics research tags for articles without a training tag.

We approached the problem as a multiclass classification task (see Hastie, Tibshirani, and Friedman, 2009), where the training data consisted of Constellate articles that already had a manually assigned field of economics research tag. To create this training data, we needed to merge the Constellate data with the field of economics research tags. However, merging the two datasets was not straightforward due to differences in DOIs between Constellate articles and the articles in the manually tagged dataset. This discrepancy arises because JSTOR and Portico commonly assign a new DOI to each article, which is different from the one originally assigned by the publishers.

To overcome this challenge, we implemented the following strategy. First, for every article i in the manually tagged dataset, we identified all Constellate articles published in the same journal and year as i. Second, we calculated the Levenshtein distance between the titles of the identified Constellate articles and article i's title.[16] Finally, we assigned the field of economics research tag from article i to the Constellate article with the lowest Levenshtein distance. As a result of this process, we obtained a training dataset consisting of 11,624 Constellate articles for which we were able to accurately assign a field of economics research tag manually.

From the training data, we created a *document-term matrix* — a matrix where each row represents an article, each column corresponds to a word/token, and the cell values indicate the frequency of each word in each document. The word frequencies were obtained from the Constellate unigram counts. Before incorporating these counts into the document-term matrix, we performed several clean-up procedures on the Constellate unigram counts. First, we converted all words to lowercase. Second, we removed leading and trailing whitespaces. Third, we eliminated non-alphanumeric characters at the beginning and end of words. We also included columns in the document-term matrix to account for words appearing in the titles of the articles.[17] To calculate word frequencies, we applied the tf-idf transformation and excluded words that did not appear in at least 0.5% of all articles, as well as words that appeared in more than 80% of all articles.[18]

We chose an L_2 regularized multinomial logistic regression as our classifier (for further details see Hastie, Tibshirani, and Friedman, 2009). This approach, employing a linear classifier with a document-term matrix as input, is commonly applied in document classification tasks (Gentzkow, Kelly, and Taddy, 2019).[19] To enhance the accuracy of the classifier, we implemented an inverse proportional weighting scheme for all training observations, taking into account the class frequencies in the input data.[20] Our final model had an inverse regularization strength equal to 1.5 (C) and a maximum number of iterations equal to 150.[21] We determined these values through a cross-validation exercise, using 80% of the data for training and 20% for validation. Table 2 presents the main performance metrics obtained in the validation set, while Table 3 shows the confusion matrix, providing additional information about the predictions in the validation set.

[16] The Levenshtein distance measures the distance between two texts as the minimum number of single-character edits required to change one word into the other. See Jurafsky and Martin (2009) for more details. Before conducting this procedure, we implemented a few clean-up procedures on both the Constellate titles and the title of article i (for example, lowercasing them, striping their leading and trailing whitespaces, removing all sequences of words inside brackets, etc.).

[17] We prefixed the string "*title_*" to each of these words. This was done to distinguish the occurrences of words in the titles of the articles from those in the bodies of the articles.

[18] The tf-idf transformation (short for term frequency-inverse document frequency transformation) is a numerical transformation in which the presence/importance given to a word in a document in a given corpus increases according to the number of times a word appears in the document and is offset by the number of documents in the corpus that contain the word. It is commonly used when representing a corpus by means of a document-term matrix. See Jurafsky and Martin (2009) for more details.

[19] As the Constellate data only provides unigram, bigram, and trigram counts and lacks full-text transcriptions, we were unable to employ more advanced approaches, such as BERT-like models (See Section 4.2).

[20] For further details, see the "balanced" option of the parameter "class_weight" in scikit-learn's logistic regression documentation (https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LogisticRegression.html).

[21] As detailed in Hastie, Tibshirani, and Friedman (2009) and Gentzkow, Kelly, and Taddy (2019), L_2 penalized models optimize a function where a regularization component is added to the loss function. Typically, this component has the following structure $\lambda \sum_{j=1}^{q} \beta_j^2$, where β_j stands for the regression coefficient associated to variable j, q stands for the total number of predictors, and λ for the regularization parameter (which is typically set by means of a cross-validation exercise). Inverse regularization strength (C) is defined as $0.01/\lambda$.

Table 2. Validation predictive performance for fields of economics research tags predicted from Constellate data

	Precision	Recall	F_1-Score	n
Applied	0.86	0.88	0.87	677
Applied theory	0.63	0.59	0.61	277
Econometric methods	0.81	0.82	0.82	190
Theory	0.94	0.94	0.94	1181
All (weighted average)	0.87	0.87	0.87	2325

Notes. Precision is equal to the proportion of all observations *predicted as belonging* to a given class that were correctly classified. Recall is equal to the proportion of all observations *effectively belonging* to a given class that were correctly classified. The F_1-Score is equal to the harmonic mean between precision and recall.

Table 3. Confusion matrix for fields of economics research tags predicted from Constellate data

	Applied	Applied theory	Econometric methods	Theory
Applied	598	59	8	12
Applied theory	62	164	12	39
Econometric methods	11	8	156	15
Theory	22	29	16	1114

Notes. Rows represent the actual class and columns represent the predicted class.

Our classifier achieves a competitive performance, having an accuracy equal to 0.87 (the trace of Table 3 divided by the sum of all its values) and a weighted average F_1-Score also equal to 0.87. As a benchmark, in Angrist et al. (2020), the authors report achieving an accuracy equal to 0.80 in a similar exercise.[22] Both Table 2 and Table 3 demonstrate the exceptional performance of our classifier in predicting applied, econometric methods, and theory tags. However, the performance declines when predicting applied theory labels. Table 2 highlights that the most frequent misclassification made by the classifier is predicting applied theory articles as applied ones.

After validating our predictive model, we proceeded to retrain it using the identical configuration as described above, but this time using the entire training set. Subsequently, we used the retrained model to predict a field of economics research tag for each Constellate article that was not included in the manually tagged dataset (n=22,999). However, for the articles within the manually tagged dataset, we preserved the manual tags and did not substitute them with the predicted ones. This yielded a total of 34,623 articles with assigned fields of economics research tags.

4.2. Assigning a field of economics research tag to all of the articles' references and citations

After assigning a field of economics research tag to each article in the Constellate corpus, we proceeded to download data on their detailed references and citations from Semantic Scholar. However, matching the data from Semantic Scholar to Constellate is not straightforward due to differences in DOIs between the two sources. To ensure a reliable merge, we took the following steps. First, we conducted a title-based search in Semantic Scholar for each of the 34,623 articles in the Constellate dataset and obtained the API response (data for all but 419 articles were successfully downloaded). Second, to ensure the accuracy of the merge, we implemented a series of filters. We excluded articles for which author information was unavailable in either Constellate or Semantic Scholar. We removed merges that did not share at least one author surname in both data sources. Additionally, we excluded merges where the same Semantic Scholar article was matched to multiple Constellate articles, ensuring that each Semantic Scholar article had a single match in Constellate. Finally, we dropped merges where the Levenshtein distance between the Constellate and Semantic Scholar titles exceeded 15. The process described above resulted in a total of 31,093 articles with detailed Constellate and Semantic

[22] An important difference between our study and Angrist et al. (2020) is that they define "styles" differently. They label what we define as applied theory as "empirical." When we recode applied theory predicted and actual tags as "applied," our classifier achieves an accuracy of 0.93 on the validation set.

Scholar data available. These articles were either referenced or cited by 2,040,265 different articles, out of which 1,330,590 (65.2%) are classified as economics research articles by Semantic Scholar.

To explore how our sample of economics research articles is cited by or references articles from different fields of economics research, we predicted a field of economics research tag for each of the 1,330,590 economics articles referenced by or citing our final sample. Once again, we treated this problem as a multiclass classification task. However, due to the absence of word counts in Semantic Scholar data, we were unable to follow the exact procedure outlined in Section 4.1. Instead, we fine-tuned a BERT-like large language model (LLM, see Devlin et al., 2019) using the articles' titles (provided by Semantic Scholar), the articles' abstracts (if available from Semantic Scholar), and the hand-labelled fields of economics research tags described in Section 3.3. The training data for this task consisted of 10,876 articles, which is a subset of the 11,624 articles used as the training set in Section 4.1. These 10,876 articles were chosen based on our high confidence in the availability and reliability of Semantic Scholar data. Specifically, we fine-tuned a pre-trained SciBERT model (Beltagy, Lo, and Cohan, 2019) to adapt it for multiclass classification. Each input text consisted of the article title (preceded by the token "*title:* ") followed by a newline character, and then followed by the article abstract (preceded by the token "*abstract:* ").[23] Training instances were weighted inversely proportional to the class frequencies in the input data. The model configuration was determined through a validation exercise, where 80% of the observations were used for training different model setups and 20% were used for validation. Table 4 presents the main performance metrics obtained in the validation set, while Table 5 shows the confusion matrix, providing additional information about the predictions in the validation set.

Table 4. Validation predictive performance for fields of economics research tags predicted from Semantic Scholar data

	Precision	Recall	F_1-Score	n
Applied	0.83	0.77	0.80	675
Applied theory	0.55	0.24	0.34	238
Econometric methods	0.76	0.62	0.68	162
Theory	0.79	0.94	0.86	1101
All (weighted average)	0.77	0.79	0.77	2176

Notes. Precision is equal to the proportion of all observations *predicted as belonging* to a given class that were correctly classified. Recall is equal to the proportion of all observations that *effectively belong* to a given class that were correctly classified. The F_1-Score is equal to the harmonic mean between precision and recall.

Table 5. Confusion matrix for fields of economics research tags predicted from Semantic Scholar data

	Applied	Applied theory	Econometric Methods	Theory
Applied	521	30	9	115
Applied theory	70	58	11	99
Econometric methods	2	2	100	58
Theory	37	16	11	1037

Notes. Rows represent the actual class and columns represent the predicted class.

Both Table 4 and Table 5 demonstrate the classifier's strong performance in predicting applied and theory labels, while achieving lower accuracy for econometric methods and applied theory articles. The overall accuracy of the classifier in the validation set is 0.79, with a weighted average F_1-Score of 0.77.[24]

Finally, we retrained the model using the same configuration as described above but with the entire training set. With this final model, we predicted a field of economics research tag for each of the 1,330,590 economics research articles referenced by or citing our sample of articles. However, as the

[23] In case of an article not having its abstract available in Semantic Scholar, the input text consisted of consisted simply of the title (preceded by the token "*title:* ").

[24] When we categorize applied theory validation labels and predictions as applied, the model achieves an accuracy of 83%. Recall that Angrist et al. (2020) achieved an accuracy of 80% in a similar exercise.

predictions in Section 4.1 demonstrated better performance than the ones presented here (a clear example of better data beating better algorithms), we retained the previous predictions for any reference or citation whose tag had already been predicted in Section 4.1.

Up to this point, our focus has been on assigning as accurate as possible fields of economics research tags to all articles published in our journals of interest, as well as to their references and the articles citing them. To achieve this, we prioritized retaining articles with manually labeled tags, even if they were not intended to be included in our final sample. After training the models as described above, we dropped two sets of articles that were not of interest to us. First, we dropped all articles published in the *Journal of Law & Economics*, the *Journal of Labor Economics*, and the *RAND Journal of Economics*, as this project centers on general research economics journals rather than field-specific ones. Second, following previous studies such as Card and DellaVigna (2013), Anauati, Galiani, and Gálvez (2016), and Anauati, Galiani, and Gálvez (2020), we filtered out all articles published in the *Papers and Proceedings* issues of *The American Economic Review*. This left a total of 24,273 economics research articles for which we collected detailed Constellate data, assigned a field of economics research tag, and collected detailed Semantic Scholar data. We refer to this set of articles as our *final sample*.[25]

4.3. Detecting latent topics in economics research articles

After constructing our final sample, our attention turns to analyzing the content of the articles. For this purpose, we employ the latent Dirichlet allocation (LDA) generative model (Blei, Ng, and Jordan, 2003). As stated in Hu (2009), LDA is an unsupervised learning algorithm that assumes words hold significant semantic information and that documents discussing similar topics tend to use similar groups of words. The algorithm automatically identifies word groups that frequently co-occur within documents in a corpus and associates them with topics. This is accomplished by modeling documents as random mixtures over latent topics, each characterized by its own distribution over the vocabulary of words. To train the LDA model, we utilized Gensim's implementation of the LDA training algorithm (see Řehůřek and Sojka, 2010). We set the number of topics to 750 and performed 20 passes through the corpus during training.[26]

LDA takes as input word counts. Before training the LDA algorithm, we conducted several preprocessing steps to clean the input text obtained from Constellate's unigram counts. These procedures aimed to enhance the quality of the training data. First, we lemmatized all tokens, converting them to their lemma form to capture the core meaning of words (e.g., replacing *walking* with *walk*, see Jurafsky and Martin, 2009). Next, we removed tokens composed solely of numbers and those not containing at least one alphanumeric character. Finally, we eliminated stop words, which are commonly used function words with little or no semantic value (e.g., *the*, *this*, etc., see Jurafsky and Martin, 2009). These steps collectively improved the quality of the tokenized text, ensuring a more effective training process for the LDA model.

Having trained the LDA algorithm on our final sample, we obtained two datasets as the main outputs. The first dataset indicates the presence of words in topics, where topics tend to be formed by a small number of words. We denote the presence of word l in topic k as $w_{k,l}$, with $w_{k,l} \geq 0$ and $\sum_l w_{k,l} = 1$. The second dataset indicates the presence of topics in documents, where documents tend to cover a small number of topics. We denote the presence of topic k in document i as $t_{i,k}$, with $t_{i,k} \geq 0$ and $\sum_k t_{i,k} = 1$. As a final step, we conducted a filtering process which excluded topics that were deemed non-meaningful (for example, a topic composed of the words *online*, *access*, *appendix*, *american*, and

[25] Recall that, for the subset of economics research articles referenced by or citing our final sample, we also assigned field of economics research tags.
[26] LDA has been widely used to uncover latent topics in research articles. Previous studies, such as Hall, Jurafsky, and Manning (2008), Amami et al. (2016), and Gálvez (2017), have successfully applied LDA in this context.

vol — which can be associated to instruction on how to access to online appendices) or that did not exhibit a clear thematic pattern (topics that were not predominantly represented by a few words). After this filtering process, we were left with a total of 695 meaningful topics. We define an article i as "containing" topic k if the presence of topic k in article i (an output of the LDA model) ranks among the top ten, out of 695, topics with the highest presence in that article. This approach ensures that each article is associated with exactly ten topics.[27]

Here are a few examples that demonstrate the patterns identified by the LDA model. After training the LDA model on our data, the algorithm identified a topic with the following five most important words: *crime* (16.54%), *criminal* (4.03%), *police* (3.80%), *arrest* (3.19%), and *rate* (2.31%) (listed in decreasing order of importance — word presence in the topic in parentheses). This topic appears to be associated with research on the economics of crime. Three articles that strongly relate to this topic are "*Why do increased arrest rates appear to reduce crime: deterrence, incapacitation, or measurement error?*" (Levitt, 1998b), "*Reexamining criminal behavior: The importance of omitted variable bias*" (Mustard, 2003) and "*Juvenile crime and punishment*" (Levitt, 1998a). Another topic identified by the LDA model has *game* (17.82%), *player* (15.30%), *payoff* (5.12%), *equilibrium* (4.65%), and *play* (4.58%) as its five most important words (listed in decreasing order of importance). This topic appears to be related to research on game theory, and three articles that have high values for this topic are "*A folk theorem for repeated sequential games*" (Wen, 2002), "*The 'folk theorem' for repeated games with complete information*" (Wen 1994), and "*Extensive form reasoning in normal form games*" (Mailath, Samuelson, and Swinkels 1993).

5. Results

We now investigate specialization trends within fields of economics research. We link specialization with patterns indicating a discipline is narrowing the topics it covers, as well as with patterns suggesting the field is receiving fewer citations from outside fields. In Section 5.1, we analyze aggregate patterns and trends within the articles comprising our main sample. In Section 5.2, we investigate how citation patterns differ based on the discipline that cites or is referenced by economics articles (e.g., what do medicine articles cite from economics? Do these patterns differ across fields of economics research?). In Section 5.3, we explore how economics articles reference other articles within the field (e.g., do applied articles tend to cite other applied articles, or do they mostly cite theory ones?).

5.1. Aggregate trends across fields of economics research

In this section, we analyze aggregate trends across fields of economics research. First, we examine the evolution in the shares of articles from different fields of economics research published over time. Then, we investigate trends in the citations received by articles across fields of economics research. Finally, we explore aggregate trends in the topics covered by our final sample of articles.

In Figure 1, we present the annual distribution of articles across different fields of economics research published in the analyzed outlets from 1970 up to and including 2016. Specifically, for each field f and year y, we plot the proportion of articles published in year y that fall under field f. We depict this distribution for the entire sample, as well as separately for the Top 5 and the non-Top 5 general research outlets considered.

[27] We tested other thresholds and consistently obtained qualitatively similar results throughout our study.

Figure 1. Trends in the yearly share of articles published by field of economics research

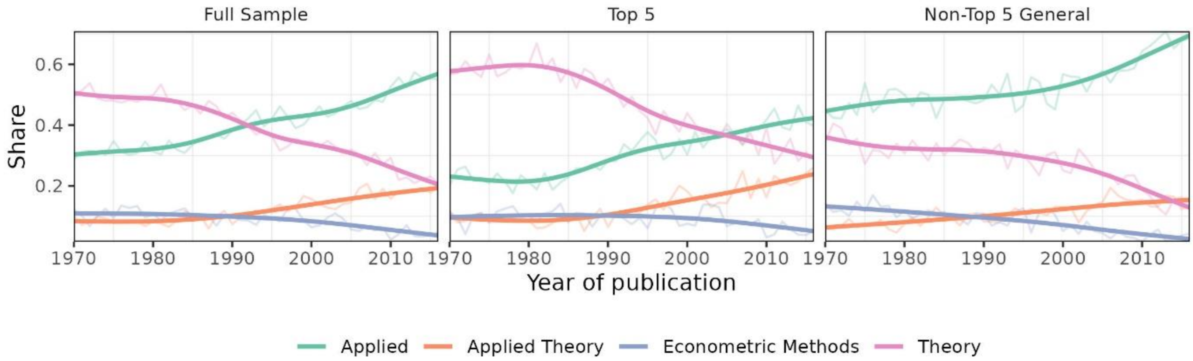

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009).

Figure 1 provides evidence of a notable shift in the publication landscape of general research journals in economics. The analysis reveals a substantial reduction in the share of theory articles, accompanied by a sharp increase in applied and applied theory articles. This finding is in line with previous studies by Hamermesh (2013), Panhans and Singleton (2017), Backhouse and Cherrier (2014), Hamermesh (2018), Angrist et al. (2017) and Angrist et al. (2020).[28] These trends are observed both in the Top 5 and non-Top 5 general research journals, although there are significant differences in the levels between the two tiers. Focusing on the Top 5 outlets, it can be seen that since 2005, applied research has become the most popular field of economics research, surpassing theory. Figure 1 also suggests that, as of the time of writing this article (late-2023), applied theory may already be the second most popular category in the Top 5 journals.[29]

Figure 2 plots trends in received citations for our sample of articles. To obtain this figure, we first calculated the sum of all citations received by articles from field f published in year y. Specifically, we considered citations received within the first seven years since publication.[30] Next, we divided this sum by the number of articles published in year y from field f, resulting in the average $\bar{c}_y^f$. The plot displays the evolution of the ratio $\bar{c}_y^f / \bar{c}_{theory}^f$ for each year y from 1970 up to and including 2016 and for all fields of economics research f — except theory, which, by definition, remains at a value of 1 for all years.[31] Values of this ratio greater than 1 for a specific field f and year y indicate that, during the first seven years since publication, papers from field f published in year y received more citations on average than theory papers published during the same period.

[28] However, it is important to note that these studies used different sample sizes, often smaller, and mostly covered shorter timeframes. Additionally, they employed different categories for fields economics of research and employed diverse tagging strategies.

[29] Recall that Constellate data is available up to 2016.

[30] Analyzing received citations in the first years since publication is a commonly adopted approach in the literature, as it has been observed that citations received during this initial period are highly correlated with citations received over longer timeframes (see, for example, Hamermesh, 2018 and Anauati, Galiani, and Gálvez, 2020).

[31] By dividing by $\bar{c}_{theory}^f$, we control for secular trends in citations, including “citation inflation.” Citation inflation refers to the phenomenon in which the number of citations received by academic papers tends to increase over time, unrelated to the intrinsic quality or impact of the research (see Neff and Olden. 2010 and Galiani and Gálvez, 2019).

Figure 2. Citations across fields of economics research relative to theory

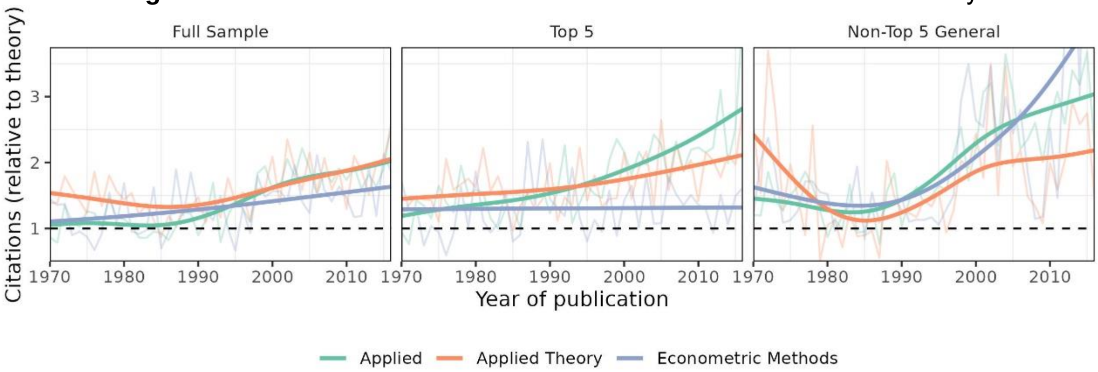

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). The dashed horizontal line at 1 serves as a visual guide, indicating the value at which the average citation of a field equals that of theory. We only consider citations received within the first seven years of publication.

Figure 2 shows that theory articles consistently had the lowest citation counts throughout the entire period (all of the plotted curves lie above 1). In the full sample, applied papers experienced a substantial increase in citations compared to theory papers, starting from the 1990s. Our data also indicate that applied theory papers received the most citations throughout the entire period, while econometric methods articles showed a steady but milder growth trend. When focusing on the Top 5 outlets, the citation ratio between econometric methods and theory articles remained steady and close to 1. In contrast, applied and applied theory articles displayed a consistently positive slope throughout the period, with a particularly strong trend observed for applied papers. It is noteworthy that Top 5 applied articles were already the most cited since the mid-1990s. Similar patterns are observed for non-Top 5 general research outlets; however, for this tier, we observe a growing gap over time between citations received by econometric methods articles relative to citations received by theory articles. Given that patterns and trends do not vary significantly between Top 5 and non-Top 5 outlets, we will only present results for our full sample of articles moving forward.

We now study overall trends in the content of articles across different fields of economics research. As explained in Section 4.3, we employed an LDA model to identify the topics present in our final sample of articles. This analysis allowed us to determine: 1) the topics covered in our corpus, 2) the presence of these topics in each article, and 3) the presence of each word in each topic. With this information, we first investigate the changes in the diversity of topics addressed within each field of economics research. To accomplish this, we adopt the approach outlined in Hall, Jurafsky, and Manning (2008). Specifically, for each field f and publication year y, we calculate topic entropy as follows:

$$H(z|f,y) = -\sum_{k=1}^{K} \hat{p}(z_k|f,y) \cdot \log(\hat{p}(z_k|f,y)),$$

where K represents the total number of topics considered (which in our case is 695), z_k denotes a specific topic, and $\hat{p}(z_k|f,y)$ represents the presence of topic z_k in articles from field f published in year y. To estimate $\hat{p}(z_k|f,y)$, we first compute the ratio between the number of articles from field f published in year y that contain topic z_k and the total number of articles from field f published in year y (denoted as $r_{z_k,f,y}$). Then, we derive $\hat{p}(z_k|f,y)$ as the quotient of $r_{z_k,f,y}$ and the sum of ratios for all topics z_j (i.e., $\sum_{j=1}^{K} r_{z_j,f,y}$).[32] High values of $H(z|f,y)$ indicate that articles from field f published in year y cover a wide range of topics, while lower values suggest a narrower focus.

[32] We divide $r_{z_k,f,y}$ by $\sum_{j=1}^{K} r_{z_j,f,y}$ to guarantee that the sum of $\hat{p}(z_k|f,y)$ across all values of k equals 1. Recall that we consider an article i as containing a particular topic if that topic is among the top ten topics with the highest presence in the article (see Section 4.3).

Figure 3. Topic entropy across fields of economics research

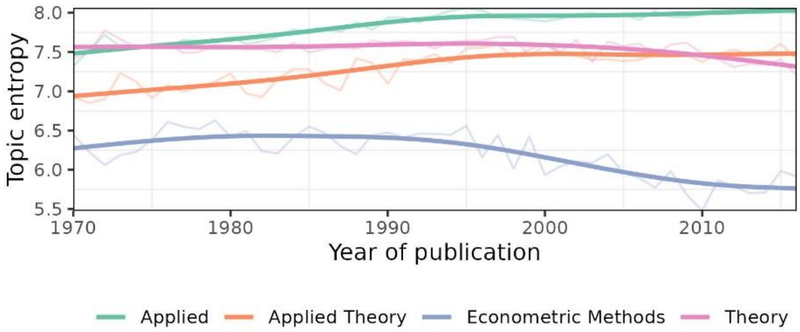

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009).

Figure 3 displays the values of $H(z|f,y)$ across different fields of economics research and publication years, revealing several noteworthy patterns. First, there are significant variations in topic entropy levels among different fields of economics research, with econometric methods exhibiting the lowest entropy. Second, trends in topic entropy differ across fields. Applied and applied theory articles demonstrated consistent positive trends throughout the entire period, with applied articles exhibiting the highest entropy as early as the 1980s. Topic entropy for theory articles remained relatively stable until the mid-1990s, after which a slight decline is observed. Econometric methods articles experienced a modest increase in topic entropy until the early 1990s, followed by a consistent decrease (reaching its lowest point by the end of the analyzed period).

To gain deeper insights into the dynamics depicted in Figure 3, we delve into the analysis of convergence and divergence in the topics covered by articles across fields of economics research. Note that Figure 3, while informative, does not distinctly reveal whether applied and applied theory articles are broadening their coverage to the point of overlap or if they cover a wider range of topics with minimal overlap between them. To quantify topic convergence/divergence, we calculated the Jensen-Shannon divergence across fields of economics research for each year in our sample. Specifically, for each pair of fields f_1 and f_2, and for every publication year y, we calculate the following expression:

$$JSD(\hat{p}(z|f_1,y) \parallel \hat{p}(z|f_2,y)) = \frac{1}{2}D(\hat{p}(z|f_1,y) \parallel M) + \frac{1}{2}D(\hat{p}(z|f_2,y) \parallel M),$$

where $\hat{p}(z|f_1,y)$ and $\hat{p}(z|f_2,y)$ are vectors representing topic presence for fields f_1 and f_2, respectively, in year y. These two vectors are obtained by concatenating the values of $\hat{p}(z_k|f_1,y)$ and $\hat{p}(z_k|f_2,y)$ for all topics k. M is equal to $(\hat{p}(z|f_1,y)+\hat{p}(z|f_2,y))/2$. $D(P \parallel Q)$ represents the Kullback-Leibler divergence calculated for P with Q as the reference.[33] Elevated values of JSD indicate increased dissimilarity or divergence between the compared probability distributions (in our case, suggesting that the fields exhibit distinct topic profiles).

In Figure 4, we illustrate trends in JSD for all pairs of fields. Each panel in the figure corresponds to a field of economics research, and within each panel, the JSD is calculated for every other field of economics research and every year considered in the study.

[33] The Kullback-Leibler divergence, denoted as $D(P \parallel Q)$, is a measure of the difference between two probability distributions. It quantifies how distribution P differs from distribution Q. It is non-symmetric and undefined when there exists an index k for which $P_k > 0$ and $Q_k = 0$. These limitations are overcome by the Jensen-Shannon divergence, which is a symmetrized version of the Kullback-Leibler divergence.

Figure 4. Jensen-Shannon divergence across fields of economics research

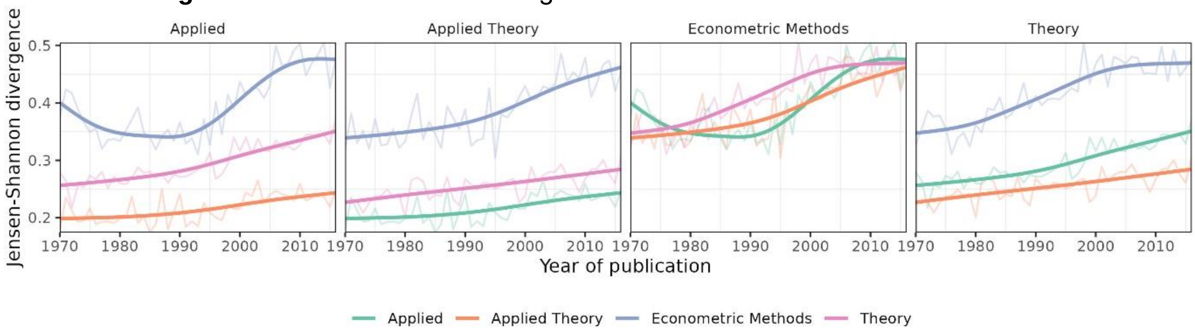

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009).

Figure 4 depicts that, irrespective of the field under consideration, the JSD between each field and any other field of economics research exhibit a consistent positive trend. This indicates that over time, all fields of economics research tended to specialize in different topics. Looking at the specific JSD levels, we observe that the field of econometric methods showed the highest topic divergence compared to any other field. On the other hand, the field of applied theory demonstrated closer topic profiles to both applied and theory fields. This finding suggests that applied theory stands between the applied and theory fields in terms of topic similarity, occupying an intermediate position. Figure 4, along with the observations from Figure 3, further suggest that both theory and econometric methods articles are becoming more specialized. In the case of applied and applied theory, the patterns suggest that while both fields are covering a wider set of topics, there is no tendency for them to overlap.

We now study which topics gained and lost prominence across fields of economics research since the mid-1990s. We identify variation in topic presence for a given field of economics research in the following way. For each field of economics research f and every topic z_k, we fit a linear regression model to the evolution of $\hat{p}(z_k|f,y)$ over y, where y ranges from 1995 up to and including 2016. This results in a set of regression slopes, where positive values indicate an increasing trend in topic presence and negative values indicate a decreasing trend. Figure 5 focuses on applied articles. The leftmost panels illustrate the trajectories of the five topics exhibiting the highest positive slopes within the applied sample, effectively highlighting the topics that gained the most prominence. Conversely, the rightmost panels depict the evolution of the five topics with the lowest (negative) slopes, representing the topics that experienced the most significant decline in prominence.

Figure 5. Topics gaining and losing more presence in applied papers since 1995

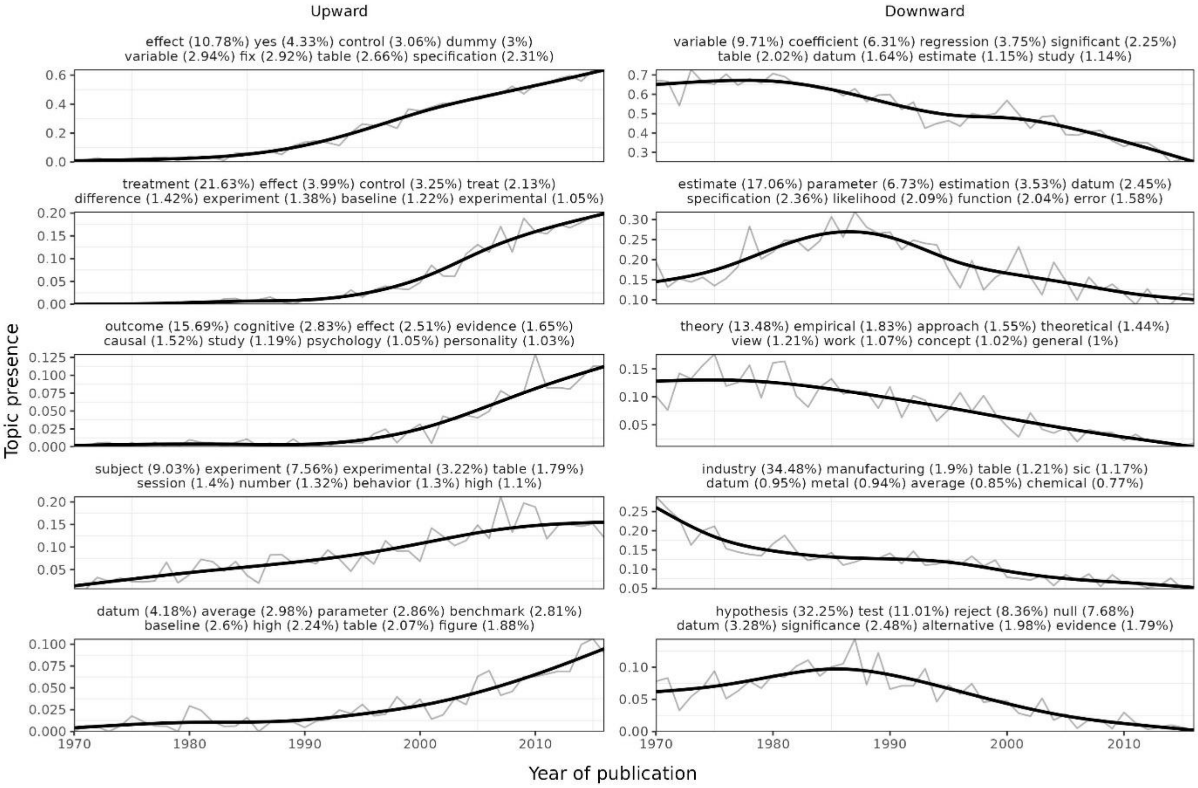

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). Each panel title displays the top 8 words that make up the corresponding topic, along with their participation within the topic (in parentheses). In the left/right column of panels, higher panels represent topics with larger/lower slope values.

Figure 5 illustrates a strong shift in the contents of applied papers. The topics showing the highest rising trends are closely related to impact and causal analysis, including words such as *effect*, *treatment*, and *control*.[34] By 2016, over 60% of all applied articles included the topic dominated by the word *effect*. On the other hand, topics associated with simple correlational analysis, indicated by words like *variable*, *coefficient*, *regression*, *table*, *specification*, and *significant*, experienced a significant decrease in their presence.[35] The figure also indicates a shift not only in statistical techniques and empirical strategies but also in domains of study. The presence of a topic related to experimental economics, including words like *subject*, *experiment*, and *experimental*, saw a substantial rise,[36] while a topic encompassing words like *industry* and *manufacture* showed a marked decline.[37]

[34] An example of a paper with a high presence of this topic is "*Income and Population Growth*" (Brueckner and Schwandt 2015).
[35] An example of a paper with a high presence of this topic is "*Income Inequality and City Size*" (Long, Rasmussen, and Haworth 1977).
[36] An example of a paper with a high presence of this topic is "*The Sequential Equilibrium Theory of Reputation Building: A Further Test*" (Neral and Ochs 1992).
[37] An example of a paper with a high presence of this topic is "*Labor Productivity and the Elasticity of Factor Substitution in West German Industries 1950-1960*" (Roskamp 1977)

Figure 6. Topics gaining and losing more presence in applied theory papers since 1995

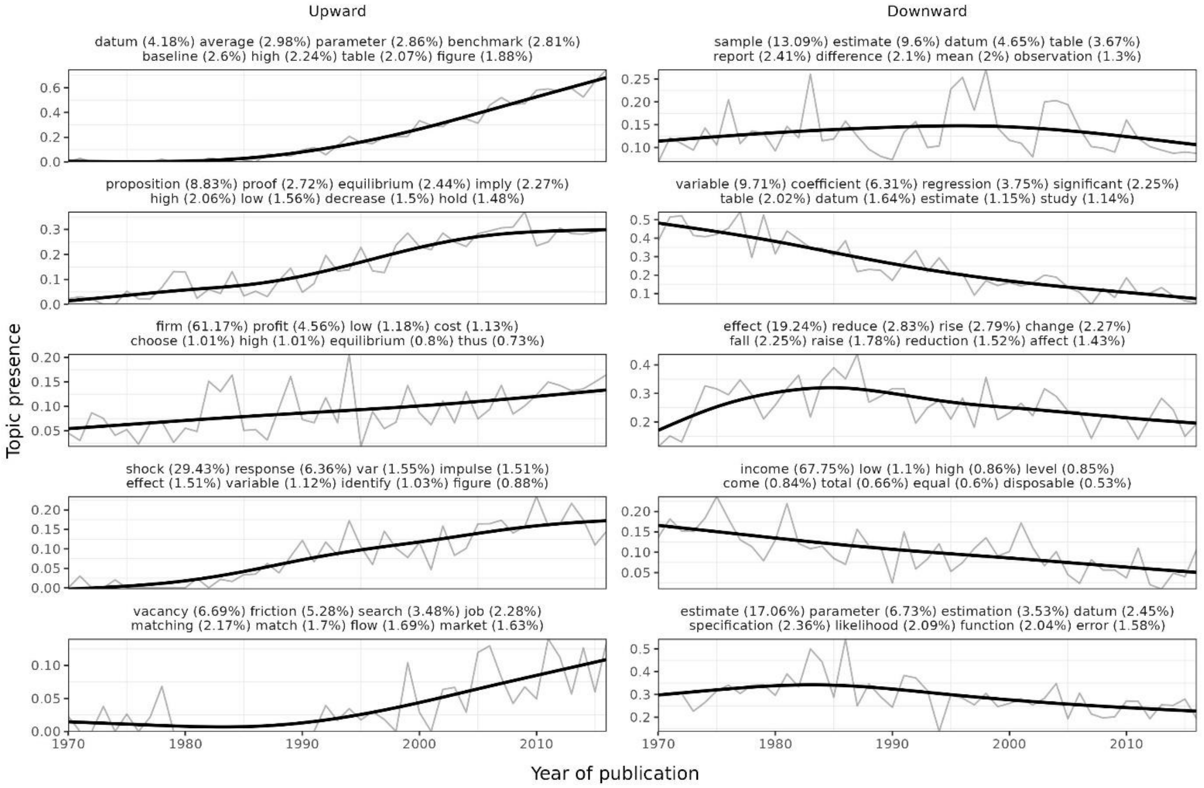

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). Each panel title displays the top 8 words that make up the corresponding topic, along with their participation within the topic (in parentheses). In the left/right column of panels, higher panels represent topics with larger/lower slope values.

Figure 6 replicates Figure 5, focusing on applied theory articles. The results show that topics commonly associated with modern macroeconomics research have significantly increased their presence since 1995. These topics include words such as *shock*, *response*, *impulse*, *friction*, and *search*.[38] Additionally, the topic centered around the word *firm* also experienced a notable rise during the same period. Interestingly, two other topics that gained prominence are 1) a topic related to data and estimation (involving words like *datum*, *average*, and *parameter*) and 2) a topic associated with theory and demonstrations (including words such as *proposition*, *proof*, and *equilibrium*). These two topics reflect the middle-ground nature of this field, combining theoretical and applied analyses. Similar to applied articles, several topics associated with simple correlational analysis showed a decline in their presence during the period.

[38] An example of a paper with a high presence of the topic lead by the word *shock* is "*Credit and Economic Activity: Credit Regimes and Nonlinear Propagation of Shocks*" (Balke 2000).

Figure 7. Topics gaining and losing more presence in econometric methods papers since 1995

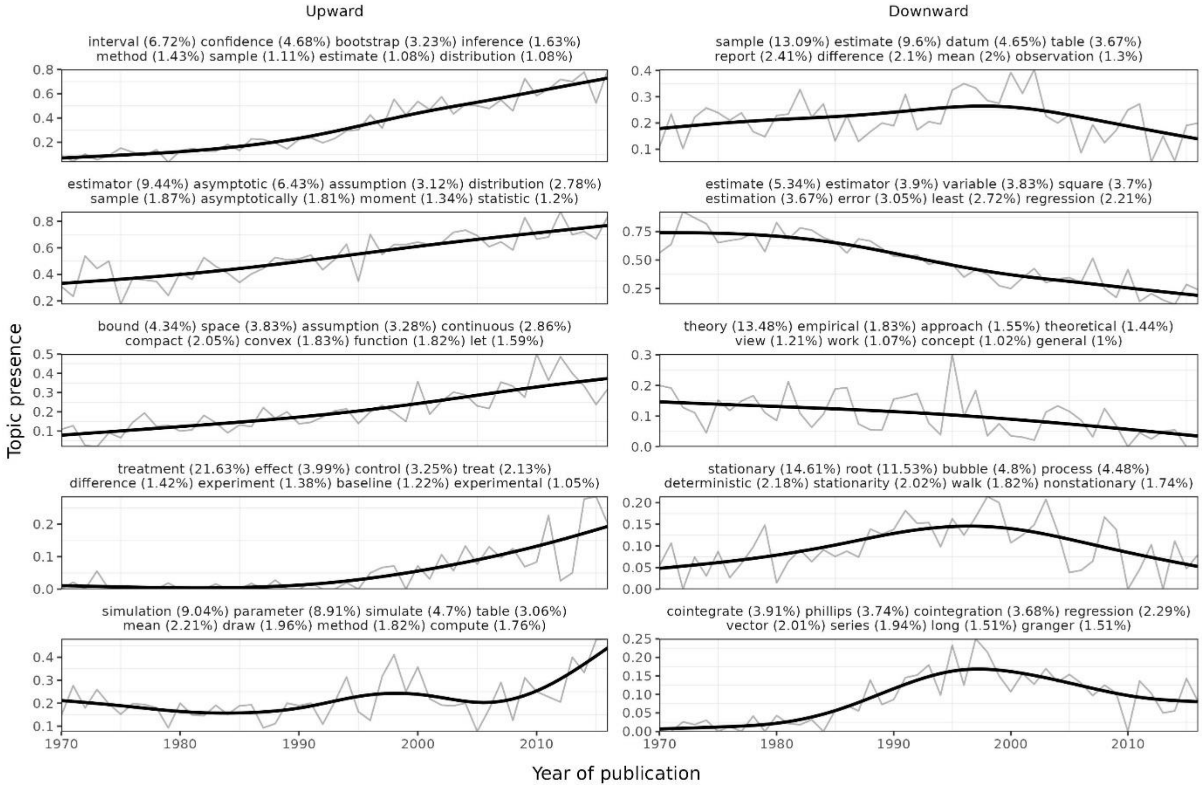

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). Each panel title displays the top 8 words that make up the corresponding topic, along with their participation within the topic (in parentheses). In the left/right column of panels, higher panels represent topics with larger/lower slope values.

Figure 7 displays the topics that have gained and lost the most prominence in econometric methods. Notably, two topics related to computational statistics, involving words like *bootstrap* and *simulation*, have experienced the most significant increase since 1995. Moreover, there has been a drastic positive trend in a topic that includes words such as *estimator* and *asymptotic*, with approximately 80% of all econometric methods articles addressing this topic by 2016 (a topic that can be readily associated with the asymptotic properties of estimators).[39] Additionally, around 50% of econometrics articles cover a topic that encompasses the words *bound*, *space*, and *assumption*, which are frequently used in the examination of the bounds of estimators. Interestingly, albeit to a lesser extent, a topic strongly associated with causal analysis (including words like *treatment*, *effect*, and *control*) has also experienced an increase over the period. Identifying clear patterns in the topics that have lost prominence since 1995 is more challenging. However, it is worth noting that two of these topics are closely related to time series analysis, with one involving words like *stationary* and *root*,[40] and the other including words such as *cointegrate*, *series*, and *granger*.[41]

[39] An example of a paper with a high presence of this topic is "*A Conditional-Heteroskedasticity-Robust Confidence Interval for the Autoregressive Parameter*" (Andrews and Guggenberger 2014).
[40] An example of a paper with a high presence of this topic is "*Pitfalls in Testing for Explosive Bubbles in Asset Prices*" (Evans 1991).
[41] An example of a paper with a high presence of this topic is "*Maximum-Likelihood Estimation of Fractional Cointegration with an Application to U.S. and Canadian Bond Rates*" (Dueker and Startz 1998).

Figure 8. Topics gaining and losing more presence in theory papers since 1995

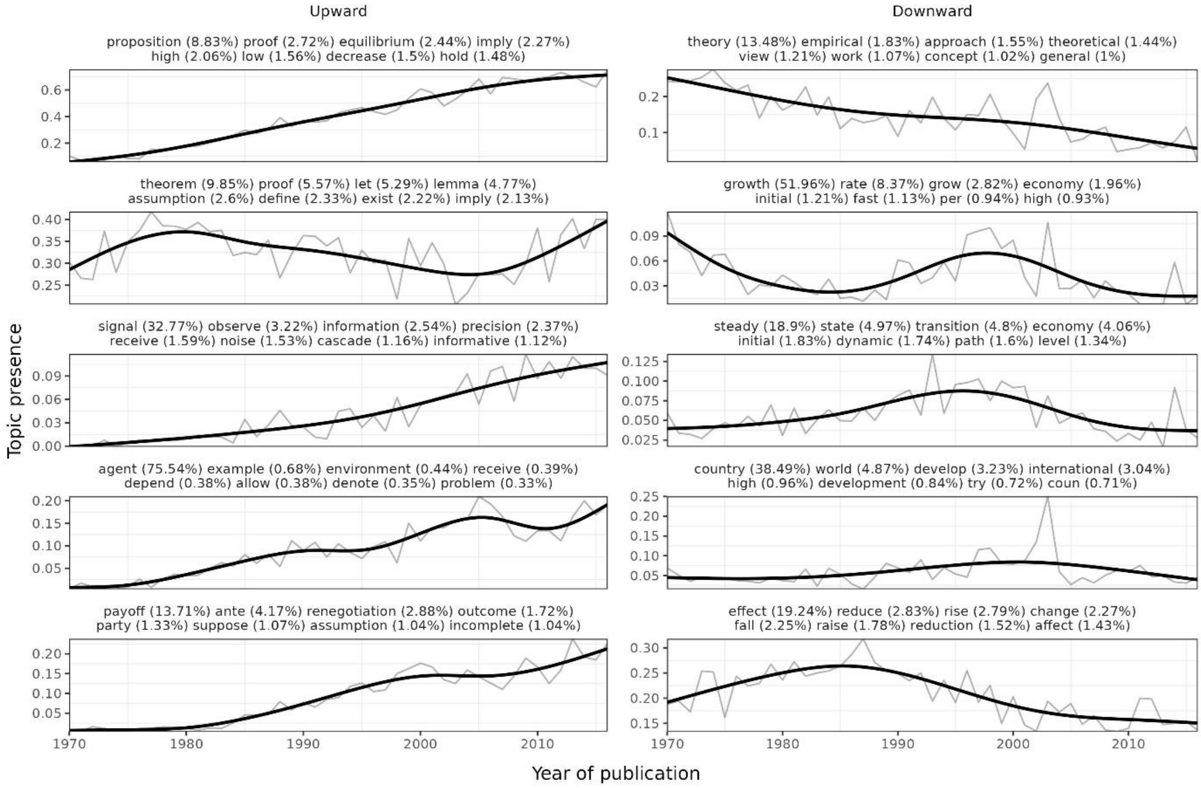

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). Each panel title displays the top 8 words that make up the corresponding topic, along with their participation within the topic (in parentheses). In the left/right column of panels, higher panels represent topics with larger/lower slope values.

Finally, Figure 8 highlights the topics that have gained and lost the most prominence since 1995 in theory articles. The two topics that have experienced the most significant increase are closely related to formal mathematical proofs, with one encompassing words like *proposition*, *proof*, and *equilibrium*,[42] and the other including words such as *theorem, proof*, *lemma*, and *assumption*.[43] Two additional topics are associated with game theory studies, with one featuring words like *signal*, *observe*, and *information*,[44] and the other including the word *payoff*.[45] A topic that primarily uses the word *agent* has also seen a rise since 1995. Notably, the topic that has experienced the sharpest decrease is predominantly composed of the word *theory*.[46] Other topics showing declining trends include a topic associated with steady-state analysis (including words like *steady*, *state*, and *path*), a topic potentially linked to growth studies (encompassing words like *growth*, *rate*, and *grow*), and a topic mainly consisting of the word *effect*.

5.2. The interplay between fields of economics research and fields of study other than economics

Up to now, we focused on studying aggregated trends across fields of economics research. In this section, we shift our focus to exploring the interactions between economics research articles and articles from other fields of study, such as medicine and psychology. Previous studies that have

[42] An example of a paper with a high presence of this topic is "*A Comparative Welfare Analysis of Electoral Systems with Endogenous Turnout*" (Kartal 2015).
[43] An example of a paper with a high presence of this topic is "*Competitive Equilibria on Turnpikes in a McKenzie Economy, II: An Asymptotic Turnpike Theorem*" (Yano 1985).
[44] An example of a paper with a high presence of this topic is "*A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades*" (Bikhchandani, Hirshleifer, and Welch 1992).
[45] An example of a paper with a high presence of this topic is "*Unforeseen Contingency and Renegotiation with Asymmetric Information*" (Lee 2008).
[46] Upon manual inspection, we observed that many of these papers emphasize the field of economics itself and use the term "economic theory" to refer to the body of knowledge generated by economics, rather than referring to the approach of studying phenomena using a theory-based strategy.

explored extramural citations include Pieters and Baumgartner (2002) and Angrist et al. (2020). Here, we extend their analysis by focusing on documenting trends consistent with specialization within economics and incorporating topic analysis.

In Figure 9, we present two panels. The leftmost panel illustrates the trends in the proportion of extramural citations received by the average article within each field of economics research. The rightmost panel shows the trends in the proportion of extramural references (citations to articles from other disciplines generated by articles in our sample) across different fields of economics research.[47]

Figure 9. Share of extramural citations and extramural references across fields of economics research

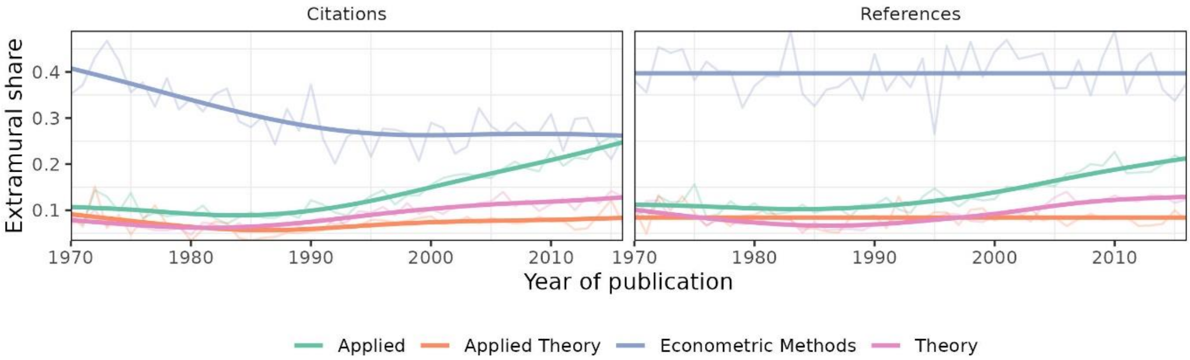

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). We only consider citations received within the first seven years since publication.

Figure 9 illustrates the variability in the extent to which economics articles are cited by and reference extramural articles across fields of economics research. Econometric methods papers receive a higher share of extramural citations and reference more extramural articles than any other field. However, while the share of extramural references has remained remarkably stable throughout the entire period, the share of extramural citations has experienced a steady decline. Applied is the field that comes second in terms of extramural citations and references. For applied papers, we observe a steady increase in both extramural citations and extramural references since the early 1990s. By 2016, approximately 25% of all citations received by applied papers were from outside economics, and nearly 20% of their references were extramural ones. Finally, both theory and applied theory papers displayed a more stable behavior during the period 1970-2016, with theory papers showing a mild but negligible upward trend, and applied theory papers showing no tendency at all.

To better understand these patterns, Table 6 presents the top twenty topics that showcase the strongest correlation with the articles' share of extramural citations. Additionally, it includes the percentage of articles containing each listed topic for each field of economics research. The calculation of these correlations is as follows. First, we determine the share of extramural citations received within the first seven years of publication for each article in our sample. Next, we calculate the Pearson correlation between the articles' shares of extramural citations and the presence of each topic across the different articles, yielding one correlation for each topic considered. Finally, we identify the topics with the highest correlation values, indicating their strong association with extramural citations.

[47] Specifically, we begin by calculating, for each article, the proportion of its citations or references that originate from or point toward fields outside of economics. For the calculation of extramural citations, we exclude articles that have not received at least one citation and only consider citations received within the first seven years of publication. Similarly, for extramural references, we exclude articles that do not reference at least one other article. Subsequently, we compute the average of these proportions for each field of economics research and publication year, and plot the results accordingly. A citation or reference is classified as extramural if the Semantic Scholar fields of study tags associated with the citing or referenced article do not include economics.

Table 6. Topics exhibiting the highest correlation with extramural citations

Topic	Correlation with the share of extramural citations	Share of articles including the topic			
		Applied	Applied Theory	Econometric Methods	Theory
estimator (9.44%) asymptotic (6.43%) assumption (3.12%) distribution (2.78%) sample (1.87%) asymptotically (1.81%) moment (1.34%) statistic (1.2%)	0.272	0.55%	1.46%	51.06%	0.41%
estimate (5.34%) estimator (3.9%) variable (3.83%) square (3.7%) estimation (3.67%) error (3.05%) least (2.72%) regression (2.21%)	0.238	7.43%	3.57%	54.26%	0.64%
matrix (25.25%) vector (19.12%) element (7.43%) diagonal (3.13%) covariance (3%) coefficient (1.5%) structure (1.35%) zero (1.35%)	0.191	0.65%	2.08%	36.27%	3.20%
outcome (15.69%) cognitive (2.83%) effect (2.51%) evidence (1.65%) causal (1.52%) study (1.19%) psychology (1.05%) personality (1.03%)	0.185	3.53%	0.66%	1.58%	1.45%
test (35.03%) power (4.77%) testing (4.35%) statistic (3.84%) alternative (2.38%) null (1.79%) sample (1.69%) regression (1.63%)	0.183	5.17%	2.81%	28.34%	0.61%
score (23.13%) grade (10.71%) test (4.2%) math (1.94%) high (1.68%) standard (1.41%) ged (1.13%) exam (1.07%)	0.173	3.18%	0.31%	1.63%	0.27%
teacher (12.92%) student (4.79%) achievement (4.29%) school (2.43%) teach (2.36%) effect (2.01%) classroom (2%) pupil (1.96%)	0.154	2.22%	0.42%	0.25%	0.27%
school (31.3%) student (4.62%) high (3.43%) education (1.67%) attend (1.66%) enrollment (1.37%) effect (1.34%) dropout (1.06%)	0.154	4.62%	1.46%	0.25%	0.48%
effect (10.78%) yes (4.33%) control (3.06%) dummy (3%) variable (2.94%) fix (2.92%) table (2.66%) specification (2.31%)	0.154	28.16%	5.27%	3.55%	0.24%
health (24.42%) illness (1.77%) medical (1.58%) mental (1.11%) age (0.96%) disease (0.85%) effect (0.84%) status (0.79%)	0.151	4.04%	1.21%	0.25%	0.47%
serial (7.4%) gmm (7.04%) moment (3.98%) autocorrelation (3.82%) correlation (2.74%) durbin (2.1%) error (1.52%) serially (1.49%)	0.146	0.85%	0.97%	12.91%	0.14%
interval (6.72%) confidence (4.68%) bootstrap (3.23%) inference (1.63%) method (1.43%) sample (1.11%) estimate (1.08%) distribution (1.08%)	0.139	1.67%	2.53%	28.29%	0.65%
student (28.71%) class (27.59%) course (6.05%) instructor (0.89%) stu (0.66%) dent (0.64%) section (0.57%) semester (0.53%)	0.138	1.91%	0.45%	0.84%	0.96%
density (31.8%) distribution (13.58%) likelihood (6.38%) maximum (3.8%) censor (3.07%) observation (2.22%) normal (1.59%) function (1.56%)	0.132	0.53%	0.76%	12.62%	0.49%
patient (22.35%) medical (9.26%) physician (8.59%) doctor (6.46%) practice (5.05%) medicine (1.62%) visit (1.38%) clinical (1.2%)	0.125	1.03%	0.35%	0.10%	0.29%
nonparametric (5.3%) estimation (5.15%) kernel (4.82%) function (4.3%) moment (4.06%) parametric (3.6%) semiparametric (2.81%) estimator (2.34%)	0.123	0.43%	1.32%	16.71%	0.31%
crime (16.54%) criminal (4.03%) police (3.8%) arrest (3.19%) rate (2.31%) offender (1.56%) violent (1.53%) property (1.39%)	0.119	1.77%	0.52%	0.25%	0.50%
treatment (21.63%) effect (3.99%) control (3.25%) treat (2.13%) difference (1.42%) experiment (1.38%) baseline (1.22%) experimental (1.05%)	0.115	6.91%	0.76%	3.55%	0.32%
sample (13.09%) estimate (9.6%) datum (4.65%) table (3.67%) report (2.41%) difference (2.1%) mean (2%) observation (1.3%)	0.115	34.70%	13.22%	22.42%	0.63%
college (26.43%) high (5.2%) student (3.77%) graduate (3.63%) admission (3.24%) tuition (2.39%) enrollment (1.79%) attend (1.75%)	0.104	1.88%	1.04%	0.10%	0.34%

Notes. Topics are sorted in decreasing order according to the estimated correlations. We only consider citations received within the first seven years of publication.

Table 6 reveals that many of the topics strongly associated with extramural citations consist of words commonly linked to fields of study such as psychology, health, criminology, and education. Interestingly, the majority of these topics are found most frequently in applied papers. Table 6 also lists topics specifically related to econometric studies (containing words such as *estimator*, *covariance*, *gmm*, and *nonparametric*), which unsurprisingly are predominantly present in econometric methods papers. Notably, neither theory nor applied theory papers exhibit the highest presence in any of the topics listed in the table.

Table 7. Topics exhibiting the lowest correlation with extramural citations

Topic	Correlation with the share of extramural citations	Share of articles including the topic			
		Applied	Applied Theory	Econometric Methods	Theory
equilibrium (43.53%) exist (1.87%) existence (1.09%) unique (0.97%) equilib (0.94%) rium (0.87%) equi (0.82%) librium (0.68%)	-0.163	1.50%	8.12%	1.23%	27.48%
price (69.94%) change (1.36%) average (0.52%) low (0.48%) high (0.38%) different (0.33%) cost (0.33%) demand (0.28%)	-0.159	14.07%	18.18%	5.22%	19.36%
effect (19.24%) reduce (2.83%) rise (2.79%) change (2.27%) fall (2.25%) raise (1.78%) reduction (1.52%) affect (1.43%)	-0.144	12.69%	24.84%	4.14%	21.78%
output (70.83%) level (3.02%) produce (2.3%) total (1.15%) function (0.95%) put (0.89%) assumption (0.7%) thus (0.54%)	-0.144	2.09%	3.26%	1.43%	3.53%
rate (64.44%) interest (11.57%) high (1.03%) expect (0.88%) level (0.79%) low (0.63%) change (0.5%) constant (0.42%)	-0.142	10.62%	13.84%	2.86%	8.59%
capital (63.92%) physical (2.19%) accumulation (2.14%) return (1.94%) stock (1.06%) investment (0.97%) economy (0.5%) neoclassical (0.49%)	-0.140	4.33%	11.73%	0.59%	7.61%
demand (54.4%) supply (22.45%) function (1.47%) aggregate (1.16%) mand (0.86%) shift (0.71%) elastic (0.55%) quantity (0.51%)	-0.138	4.29%	5.83%	2.86%	6.92%
real (49.69%) nominal (13.39%) rate (1.62%) interest (1.43%) level (0.72%) price (0.7%) anticipate (0.7%) effect (0.69%)	-0.133	3.58%	5.17%	1.08%	3.49%
consumption (54.89%) intertemporal (3.64%) consume (3.06%) sumption (2.08%) consump (1.98%) con (1.82%) tion (0.8%) marginal (0.7%)	-0.133	3.28%	9.23%	0.54%	4.56%
economy (28.96%) pareto (7.79%) competitive (2.89%) feasible (1.55%) initial (1.36%) endowment (1.29%) every (1.27%) equilibrium (1.18%)	-0.120	0.81%	2.57%	0.79%	10.37%
period (63.73%) begin (2.07%) end (1.77%) time (1.12%) next (1.09%) future (1.01%) previous (0.95%) current (0.93%)	-0.109	3.27%	4.61%	1.77%	6.78%
market (76.39%) structure (0.97%) mar (0.64%) ket (0.52%) effect (0.37%) clearing (0.29%) kets (0.29%) different (0.28%)	-0.108	6.08%	5.66%	1.03%	6.96%
money (61.72%) monetary (2.65%) interest (2.59%) friedman (2.41%) holding (2.22%) velocity (2.1%) hold (1.8%) level (1.09%)	-0.108	2.99%	2.64%	0.94%	4.49%
aggregate (44.95%) tfp (5.03%) across (1.56%) aggregation (1.31%) aggre (1.19%) gate (1.14%) decomposition (1.06%) average (0.96%)	-0.108	1.47%	3.40%	1.28%	0.87%
labor (57.82%) supply (7.62%) force (4.97%) wage (1.86%) work (1.56%) rate (0.84%) income (0.65%) market (0.63%)	-0.107	5.41%	8.85%	1.13%	3.47%
good (63.9%) produce (2.49%) consume (1.32%) bad (0.74%) thus (0.55%) must (0.54%) different (0.42%) let (0.38%)	-0.105	1.63%	4.51%	0.49%	8.57%
commodity (31.67%) producer (15.63%) price (5.77%) com (1.29%) supply (1.14%) modity (0.9%) kemp (0.72%) stabilization (0.71%)	-0.104	1.15%	2.32%	0.64%	4.59%
utility (41.65%) function (4.62%) maximize (1.97%) individual (1.01%) maximization (1.01%) concave (0.9%) preference (0.83%) taste (0.81%)	-0.103	1.52%	5.10%	2.61%	13.41%
wage (70.29%) high (1.18%) low (0.74%) hourly (0.67%) level (0.61%) worker (0.5%) determination (0.48%) labor (0.47%)	-0.101	8.05%	10.20%	1.18%	5.64%
taxis (9.27%) taxation (7.86%) distortion (7.64%) tax (5.09%) lump (3.99%) ramsey (2.21%) distort (1.77%) public (1.57%)	-0.101	0.64%	4.34%	0.30%	3.97%

Notes. Topics are sorted in increasing order according to the estimated correlations. We only consider citations received within the first seven years of publication.

Table 7 replicates Table 6, but presents the twenty topics that exhibit the strongest negative correlation with the articles' share of extramural citations. It illustrates that topics with lower correlations with extramural citations predominantly consist of words categorized as economics jargon. These topics are primarily found in applied theory and theory papers. Notably, there are no instances where applied or econometric methods papers exhibit the highest topic presence among the listed topics.

Figure 10 visualizes how different fields of study interact with various fields of economics research. This analysis was conducted as follows: first, for each article i in our sample that received at least one extramural citation (within its first seven years since publication), we calculated the proportion of article i's extramural citations received within the first seven years since publication that originate from each field of study. Then, having calculated these shares, we averaged these values over all articles published in a given year and belonging to a given field of economics research. Figure 10 displays these averages across fields of study, fields of economics research, and years of publication.

Figure 10. Origin of extramural citations by year of publication and field of economics research

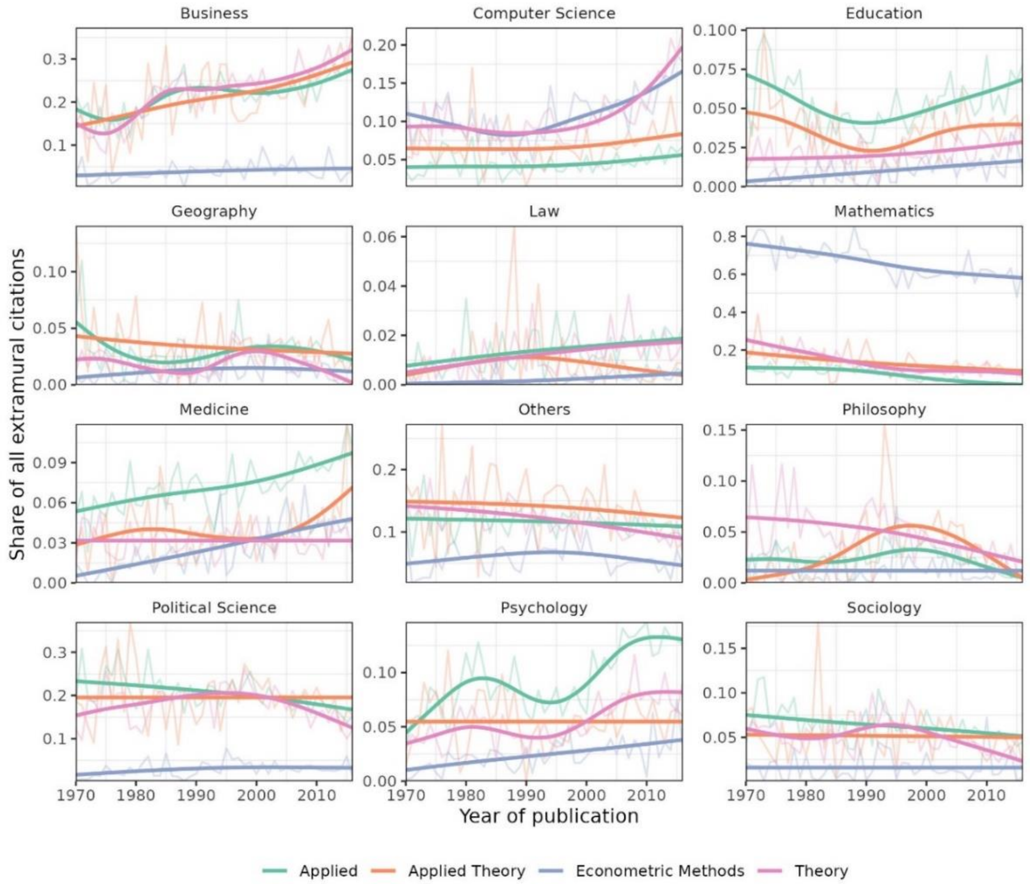

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). “Others” includes: agricultural and food sciences, art, biology, chemistry, engineering, environmental science, geology, history, linguistics, materials science, and physics. We only consider citations received within the first seven years of publication.

Figure 10 shows several interesting patterns in the origin of extramural citations. First, the origin of citations varies greatly across fields of economics research. For example, mathematics represents a large share of econometric methods citations. Computer science represents a large share of both theory and econometric methods citations (a pattern also documented in Angrist et al., 2020). Business is the field of study that cites applied, applied theory, and theory the most, almost not citing econometric methods at all. Second, trends differ across fields of economics research and disciplines. For instance, there has been a significant increase in the number of citations from the field of medicine for applied papers. This trend is not observed in theory articles. By 2016, approximately 10% of all applied papers’ extramural citations were sourced from medicine. Conversely, theory articles demonstrate a marked upward trend in citations from computer science. This pattern is also seen in econometric methods articles. Third, certain fields of study appear to have diminished in terms of their importance in citations to economics papers. This decline is evident in political science, philosophy, geography, sociology, and the disciplines categorized under "Others" (when considered as a whole).

To gain further insights into the patterns depicted in Figure 10, we present Table 8, which showcases the topics exhibiting the strongest correlation with extramural citations from each field of study. The methodology employed in constructing this table is akin to that used for Table 6, with the difference that it considers the share of citations from a specific field of study, rather than the share of received extramural citations relative to all received citations. Table 8 also provides, by field of economics research, the proportion of papers featuring each topic.

Table 8. Topics associated with citations from various fields of study

Field of study	Topic	Correlation with citations shares	Share of articles including the topic			
			Applied	Applied Theory	Econometric Methods	Theory
Business	firm (36.41%) datum (1.9%) average (1.1%) sample (1.07%) table (0.99%) number (0.86%) measure (0.85%) total (0.81%)	0.335	7.74%	4.96%	0.39%	0.38%
	performance (26.22%) manager (22.74%) management (7.71%) managerial (4.95%) practice (1.56%) manage (1.55%) perform (1.45%) mance (1.23%)	0.231	1.75%	0.90%	0.34%	1.62%
	shareholder (7.41%) takeover (2.8%) governance (2.11%) control (1.79%) tender (1.63%) journal (1.28%) corporate (1.17%) shleifer (1.03%)	0.227	1.10%	0.73%	0.05%	1.85%
Computer Science	game (17.82%) player (15.3%) payoff (5.12%) equilibrium (4.65%) play (4.58%) strategy (1.63%) outcome (1.23%) subgame (0.93%)	0.239	3.02%	1.80%	0.54%	9.92%
	step (22.89%) algorithm (15.92%) number (3.76%) shapley (3.51%) grid (2.65%) visit (2.4%) compute (2.16%) find (2.01%)	0.235	0.26%	0.31%	3.45%	1.31%
	theorem (9.85%) proof (5.57%) let (5.29%) lemma (4.77%) assumption (2.6%) define (2.33%) exist (2.22%) imply (2.13%)	0.195	0.39%	4.16%	43.62%	33.49%
Education	school (31.3%) student (4.62%) high (3.43%) education (1.67%) attend (1.66%) enrollment (1.37%) effect (1.34%) dropout (1.06%)	0.503	4.62%	1.46%	0.25%	0.48%
	teacher (12.92%) student (4.79%) achievement (4.29%) school (2.43%) teach (2.36%) effect (2.01%) classroom (2%) pupil (1.96%)	0.463	2.22%	0.42%	0.25%	0.27%
	college (26.43%) high (5.2%) student (3.77%) graduate (3.63%) admission (3.24%) tuition (2.39%) enrollment (1.79%) attend (1.75%)	0.431	1.88%	1.04%	0.10%	0.34%
Geography	city (46.55%) urban (3.87%) population (1.45%) smsa (1.11%) live (0.9%) york (0.86%) metropolitan (0.82%) across (0.76%)	0.219	2.15%	1.42%	0.00%	0.41%
	location (23.88%) spatial (8.68%) locate (5.42%) center (2.99%) agglomeration (2.93%) distance (2.15%) commute (1.81%) geographic (1.46%)	0.178	1.42%	1.42%	0.79%	1.26%
	migration (29.62%) migrant (13.25%) origin (5.58%) migrate (4.68%) destination (4.2%) emigration (2.08%) united (1.35%) flow (1.21%)	0.151	1.31%	0.31%	0.05%	0.33%
Law	crime (16.54%) criminal (4.03%) police (3.8%) arrest (3.19%) rate (2.31%) offender (1.56%) violent (1.53%) property (1.39%)	0.588	1.77%	0.52%	0.25%	0.50%
	court (10.55%) judge (8.34%) sentence (4%) defendant (2.88%) judicial (2.5%) justice (2.28%) incarceration (2.01%) convict (1.92%)	0.458	0.81%	0.14%	0.10%	0.58%
	etal (13.38%) settlement (11.93%) award (7.18%) dispute (4.04%) plaintiff (3.82%) litigation (3.64%) settle (2.95%) suit (2.11%)	0.217	0.52%	0.28%	0.05%	0.40%
Mathematics	estimator (9.44%) asymptotic (6.43%) assumption (3.12%) distribution (2.78%) sample (1.87%) asymptotically (1.81%) moment (1.34%) statistic (1.2%)	0.547	0.55%	1.46%	51.06%	0.41%
	matrix (25.25%) vector (19.12%) element (7.43%) diagonal (3.13%) covariance (3%) coefficient (1.5%) structure (1.35%) zero (1.35%)	0.424	0.65%	2.08%	36.27%	3.20%
	estimate (5.34%) estimator (3.9%) variable (3.83%) square (3.7%) estimation (3.67%) error (3.05%) least (2.72%) regression (2.21%)	0.422	7.43%	3.57%	54.26%	0.64%
Medicine	health (24.42%) illness (1.77%) medical (1.58%) mental (1.11%) age (0.96%) disease (0.85%) effect (0.84%) status (0.79%)	0.476	4.04%	1.21%	0.25%	0.47%
	patient (22.35%) medical (9.26%) physician (8.59%) doctor (6.46%) practice (5.05%) medicine (1.62%) visit (1.38%) clinical (1.2%)	0.394	1.03%	0.35%	0.10%	0.29%
	hospital (22.71%) patient (3.22%) admission (2.82%) medicare (2.3%) hmo (1.65%) bed (1.63%) health (1.6%) roo (1.04%)	0.387	1.22%	0.31%	0.20%	0.21%
Others	century (5.42%) historical (3.82%) history (3.41%) press (2.26%) early (1.83%) nineteenth (1.73%) modern (1.49%) cambridge (1.2%)	0.062	0.32%	0.13%	0.00%	0.28%
	pollution (11.78%) environmental (11%) emission (9.1%) air (6.22%) abatement (2.42%) permit (2.2%) clean (2.14%) level (1.92%)	0.055	0.33%	0.41%	0.04%	0.24%
	electricity (8.22%) power (5.42%) electric (4.4%) gas (2.57%) energy (1.81%) generator (1.8%) load (1.71%) generation (1.61%)	0.051	0.48%	0.19%	0.09%	0.06%
Philosophy	theory (13.48%) empirical (1.83%) approach (1.55%) theoretical (1.44%) view (1.21%) work (1.07%) concept (1.02%) general (1%)	0.120	7.25%	6.38%	10.94%	14.94%
	religious (11.57%) religion (8.12%) church (7.79%) catholic (6.41%) protestant (3.41%) sacrifice (2.33%) catholics (1.48%) secular (1.45%)	0.062	0.61%	0.24%	0.05%	0.11%
	welfare (33.92%) social (15.68%) alternative (1.75%) pareto (1.31%) socially (1.19%) function (0.92%) choice (0.85%) sen (0.66%)	0.059	1.15%	4.72%	1.18%	7.24%
Political Science	political (15.51%) power (6.11%) politician (3.69%) citizen (3.05%) public (1.75%) politic (1.5%) democracy (0.96%) persson (0.9%)	0.303	2.34%	2.39%	0.15%	3.62%
	vote (26.3%) voter (13.19%) outcome (1.55%) turnout (1.42%) election (1.34%) issue (1.06%) number (0.97%) pivotal (0.94%)	0.246	2.06%	2.12%	0.10%	3.19%
	election (17.22%) party (6.31%) electoral (6.26%) political (3.89%) incumbent (2.57%) presidential (2.56%) partisan (1.57%) reelection (1.38%)	0.222	1.60%	1.25%	0.10%	0.64%
Psychology	outcome (15.69%) cognitive (2.83%) effect (2.51%) evidence (1.65%) causal (1.52%) study (1.19%) psychology (1.05%) personality (1.03%)	0.375	3.53%	0.66%	1.58%	1.45%
	mother (10.63%) child (10.3%) effect (2.67%) father (1.39%) maternal (1.35%) age (1.18%) sample (0.99%) time (0.85%)	0.291	2.58%	1.01%	0.25%	0.13%
	subject (9.03%) experiment (7.56%) experimental (3.22%) table (1.79%) session (1.4%) number (1.32%) behavior (1.3%) high (1.1%)	0.290	9.41%	1.98%	3.84%	1.37%
Sociology	woman (33.63%) man (22.15%) work (1.81%) marry (1.64%) force (1.42%) difference (1.3%) less (1%) single (0.93%)	0.182	4.07%	1.73%	0.44%	0.45%
	religious (11.57%) religion (8.12%) church (7.79%) catholic (6.41%) protestant (3.41%) sacrifice (2.33%) catholics (1.48%) secular (1.45%)	0.169	0.61%	0.24%	0.05%	0.11%
	crime (16.54%) criminal (4.03%) police (3.8%) arrest (3.19%) rate (2.31%) offender (1.56%) violent (1.53%) property (1.39%)	0.166	1.77%	0.52%	0.25%	0.50%

Notes. "Others" includes: agricultural and food sciences, art, biology, chemistry, engineering, environmental science, geology, history, linguistics, materials science, and physics. Within each citing field of study, topics are sorted in decreasing order according to the estimated correlations. We only consider citations received within the first seven years of publication.

Table 8 shows that different fields of study tend to cite articles containing different topics. The observed patterns align with expectations. For example, education articles tend to cite economics articles featuring topics containing the words *school*, *teacher*, and *college*,[48] medicine articles tend to cite economics articles with topics containing words such as *health*, *patient*, and *hospital*,[49] and computer science articles tend to cite economics articles with topics centered around words such as *game*, *algorithm*, and *theorem*.

Finally, Figure 11 examines the citation behavior of different fields of economics research towards other disciplines. It presents an equivalent analysis to Figure 10 but focuses on the share of extramural references to specific fields of study.

Figure 11. Destination of extramural references by year of publication and field of economics research

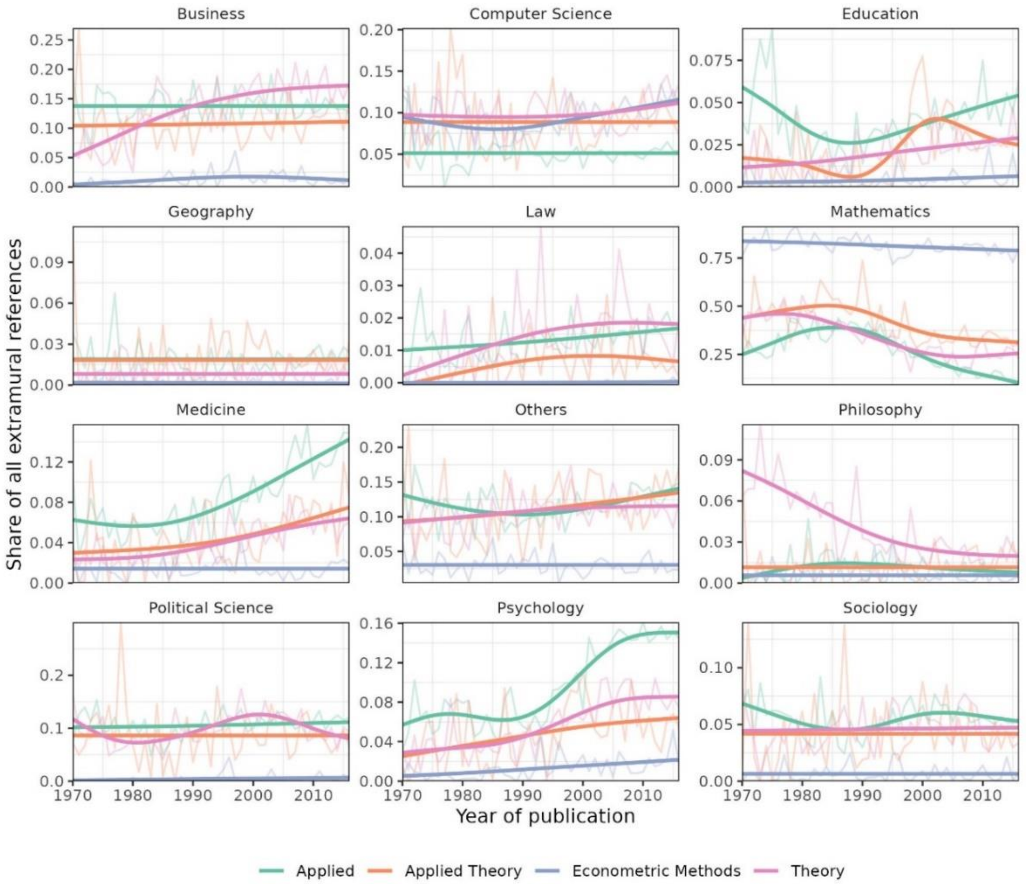

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). "Others" includes: agricultural and food sciences, art, biology, chemistry, engineering, environmental science, geology, history, linguistics, materials science, and physics.

Figure 11 reveals a compelling resemblance to Figure 10, as the patterns observed in both figures align quite closely. This parallelism suggests a strong interaction between some fields of study other than economics and different fields of economics research.

[48] An example of a paper with a high presence of this topic is "*Do Students Benefit from Attending Better Schools? Evidence from Rule-based Student Assignments in Trinidad and Tobago*" (Jackson 2010).
[49] An example of a paper with a high presence of this topic is "*Socioeconomic Status and Child Health: Why Is the Relationship Stronger for Older Children?*" (Currie and Stabile 2003).

5.3. The interplay within and between fields of economics research

In this final section, our focus shifts to examining the dynamics of citations within and between fields of economics research. To achieve this, we utilize the field of economics research tags assigned to each economics research article citing our sample of articles or being referenced by it (see Section 4.2).

The upper panels of Figure 12 depict, for every field of economics research, the relative share of citations from other economics articles that come from each field of economics research. For instance, by 2016, nearly 80% of all citations from economics papers to applied economics papers originated from other applied papers. The lower panels of Figure 12 illustrate, for every field of economics research, the relative share of economics articles cited by our sample of articles that belong to each field of economics research. For example, by 2016, almost 70% of all references in our sample of applied papers that cited other economics articles were to other applied papers.[50]

Figure 12. Citations and references within and between fields of economics research

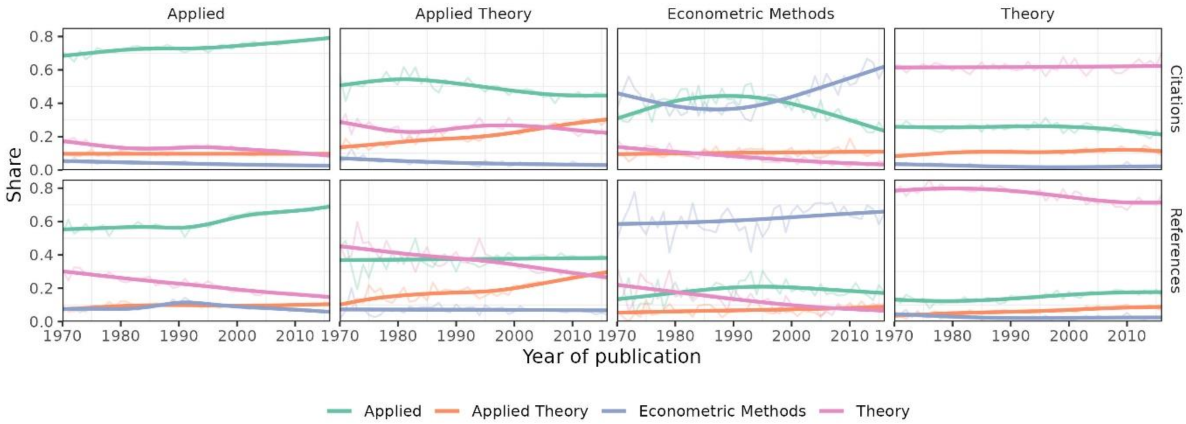

Notes. Trends are smoothed by fitting a generative additive model (GAM) to the data. See Hastie, Tibshirani, and Friedman (2009). We only consider citations received within the first seven years of publication.

Figure 12 illustrates a clear tendency towards homophily in citation behavior. This is evident from both the upper and lower panels. Homophily indicates that articles within a specific field predominantly reference and are cited by other articles within the same field. For instance, in 2016, approximately 60% of all citations received by theory articles from economics articles originated from other theory papers, and about 70% of their references to other economics articles were directed towards other theory papers. The only evident exception to this pattern is observed for applied theory articles, which demonstrate significant interactions with both applied and theory articles. Analyzing the four bottom panels collectively reveals that, by 2016, applied papers were consistently among the most referenced articles across all fields examined. Regardless of the field under analysis, applied ranked either as the most referenced or the second most referenced category. When examining trends over time, it becomes apparent that the inclination of articles to be cited by or to reference articles from their own field has intensified over the years, except in the case of theory. For theory articles, a relatively stable trend is observed in terms of citations, while a slight declining trend is observed in terms of references.

Finally, we study, for every field of economics research, which topics are cited the most by articles from different fields of economics research. For example, which are the topics present in econometric methods papers that applied papers cite most? Which are the topics present in applied

[50] The construction of this figure closely follows the methodology described for Figure 10 and Figure 11, with two key distinctions. First, it focuses on citations to and from articles within the field of economics. Second, it calculates averages across different fields of economics research instead of fields of study outside of economics. Similar to the previous figures, we only consider citations received within the first seven years since publication.

papers that theory papers cite most? Etc. We begin by studying citations to applied papers. Table 9 lists, for our sample of applied papers, which are the topics whose presence correlates the most with citations coming from different fields of economics research.[51] It also presents the proportion of applied papers containing the topic. Only topics with a presence of at least 5% are included in the table.

Table 9. Topics cited the most in applied papers by articles from different fields of economics research

Citing field of economics research	Topic	Correlation with citations shares	Share of articles including the topic
Applied	effect (10.78%) yes (4.33%) control (3.06%) dummy (3%) variable (2.94%) fix (2.92%) table (2.66%) specification (2.31%)	0.206	27.72%
	sample (13.09%) estimate (9.6%) datum (4.65%) table (3.67%) report (2.41%) difference (2.1%) mean (2%) observation (1.3%)	0.126	34.47%
	age (42.35%) old (10.08%) young (8.81%) year (1.82%) effect (0.87%) datum (0.65%) decline (0.63%) rate (0.63%)	0.092	5.06%
	decline (5.67%) change (4.57%) rise (4.08%) fall (1.81%) figure (1.4%) high (1.4%) within (1.29%) percentile (1.21%)	0.083	7.14%
	treatment (21.63%) effect (3.99%) control (3.25%) treat (2.13%) difference (1.42%) experiment (1.38%) baseline (1.22%) experimental (1.05%)	0.070	6.79%
	survey (17.94%) response (8.91%) respondent (8.35%) question (8.18%) ask (4.53%) answer (3.01%) report (2.17%) wtp (1.33%)	0.066	5.27%
Applied Theory	estimate (17.06%) parameter (6.73%) estimation (3.53%) datum (2.45%) specification (2.36%) likelihood (2.09%) function (2.04%) error (1.58%)	0.129	17.81%
	price (69.94%) change (1.36%) average (0.52%) low (0.48%) high (0.38%) different (0.33%) cost (0.33%) demand (0.28%)	0.068	14.14%
	equation (54.73%) variable (8.14%) exogenous (1.92%) determine (1.38%) endogenous (1.32%) form (1.31%) equa (0.95%) substitute (0.92%)	0.065	6.21%
	cost (70.53%) low (1.34%) fix (1.33%) incur (1.03%) total (0.95%) costly (0.88%) average (0.86%) high (0.65%)	0.059	5.80%
	datum (11.65%) quarterly (7.57%) quarter (7.56%) gap (5.57%) estimate (3.95%) table (2.1%) revision (1.2%) available (1.15%)	0.043	6.14%
	time (31.07%) constant (2.28%) zero (1.37%) equal (1.14%) initial (1.03%) assumption (1.03%) section (0.96%) interval (0.9%)	0.039	5.12%
Econometric Methods	test (35.03%) power (4.77%) testing (4.35%) statistic (3.84%) alternative (2.38%) null (1.79%) sample (1.69%) regression (1.63%)	0.168	5.17%
	estimate (5.34%) estimator (3.9%) variable (3.83%) square (3.7%) estimation (3.67%) error (3.05%) least (2.72%) regression (2.21%)	0.138	7.47%
	estimate (17.06%) parameter (6.73%) estimation (3.53%) datum (2.45%) specification (2.36%) likelihood (2.09%) function (2.04%) error (1.58%)	0.119	17.81%
	equation (54.73%) variable (8.14%) exogenous (1.92%) determine (1.38%) endogenous (1.32%) form (1.31%) equa (0.95%) substitute (0.92%)	0.111	6.21%
	hypothesis (32.25%) test (11.01%) reject (8.36%) null (7.68%) datum (3.28%) significance (2.48%) alternative (1.98%) evidence (1.79%)	0.101	5.13%
	datum (11.65%) quarterly (7.57%) quarter (7.56%) gap (5.57%) estimate (3.95%) table (2.1%) revision (1.2%) available (1.15%)	0.085	6.14%
Theory	theory (13.48%) empirical (1.83%) approach (1.55%) theoretical (1.44%) view (1.21%) work (1.07%) concept (1.02%) general (1%)	0.130	7.43%
	subject (9.03%) experiment (7.56%) experimental (3.22%) table (1.79%) session (1.4%) number (1.32%) behavior (1.3%) high (1.1%)	0.094	9.34%
	cost (70.53%) low (1.34%) fix (1.33%) incur (1.03%) total (0.95%) costly (0.88%) average (0.86%) high (0.65%)	0.079	5.80%
	price (69.94%) change (1.36%) average (0.52%) low (0.48%) high (0.38%) different (0.33%) cost (0.33%) demand (0.28%)	0.048	14.14%
	market (76.39%) structure (0.97%) mar (0.64%) ket (0.52%) effect (0.37%) clearing (0.29%) kets (0.29%) different (0.28%)	0.036	6.10%
	gold (4.2%) british (3.99%) britain (3.46%) united (2.78%) london (2.48%) england (2.02%) kingdom (1.72%) great (1.25%)	0.035	5.95%

Notes. Within each citing field of economics research, topics are sorted in decreasing order according to the estimated correlations. Only topics with a presence of at least 5% in all applied papers are listed. We only consider citations received within the first seven years of publication.

[51] We construct Table 9 using a similar approach to Table 6. Here, the correlations are computed for our applied papers sample, focusing on citations received from each field of economics research. In this way, for each topic, we calculate four correlations. We limit our consideration to citations received within seven years of publication. Only topics with a presence of at least 5% are included in the table.

Table 9 shows that the topics present in our sample of applied papers that are cited the most vary greatly depending on the field of the citing articles. Applied papers citing other applied papers tend to cite articles that include topics commonly associated with causal analysis (composed of words such as *effect* and *treatment*) and data description/construction (including words such as *sample* and *survey*). Applied theory articles citing applied papers tend to cite papers including jargon commonly associated with the estimation of models' parameters (*estimate*, *parameter*, *equation*, *variable*, etc.), but they also cite a topic led by the word *price* and another led by the word *cost*. Econometric methods mostly cite applied papers which include topics associated with estimators (e.g., a topic led by the words *test* and *power*, and another topic led by the words *estimate* and *estimator* are listed). Theory papers tend to cite applied articles containing a topic led by the word *theory*, and, as it was the case for citations coming from applied theory papers, they also cite a topic led by the word *price* and another led by the word *cost*. Interestingly, theory papers also tend to cite applied articles containing the topic led by the word *subject* and *experiment*, suggesting that applied experimental economics is an interesting research topic for this field.

Table 10, Table 11, and Table 12 replicate Table 9, but focus on citations to our sample of applied theory, econometric methods, and theory articles, respectively. A consistent observation across all tables is that the topics most frequently cited in articles from a particular field of economics research exhibit substantial variation depending on the field from which the citations originate. This implies that the topics considered important or interesting within a specific field of economics research differ greatly based on the field of the citing article.

Table 10. Topics cited the most in applied theory papers by articles from different fields of economics research

Citing field of economics research	Topic	Correlation with citations shares	Share of articles including the topic
Applied	effect (10.78%) yes (4.33%) control (3.06%) dummy (3%) variable (2.94%) fix (2.92%) table (2.66%) specification (2.31%)	0.172	5.22%
	variable (9.71%) coefficient (6.31%) regression (3.75%) significant (2.25%) table (2.02%) datum (1.64%) estimate (1.15%) study (1.14%)	0.159	21.59%
	sample (13.09%) estimate (9.6%) datum (4.65%) table (3.67%) report (2.41%) difference (2.1%) mean (2%) observation (1.3%)	0.143	13.03%
	estimate (17.06%) parameter (6.73%) estimation (3.53%) datum (2.45%) specification (2.36%) likelihood (2.09%) function (2.04%) error (1.58%)	0.133	28.44%
	decline (5.67%) change (4.57%) rise (4.08%) fall (1.81%) figure (1.4%) high (1.4%) within (1.29%) percentile (1.21%)	0.077	5.08%
	industry (34.48%) manufacturing (1.9%) table (1.21%) sic (1.17%) datum (0.95%) metal (0.94%) average (0.85%) chemical (0.77%)	0.063	6.86%
Applied Theory	datum (4.18%) average (2.98%) parameter (2.86%) benchmark (2.81%) baseline (2.6%) high (2.24%) table (2.07%) figure (1.88%)	0.305	27.90%
	shock (29.43%) response (6.36%) var (1.55%) impulse (1.51%) effect (1.51%) variable (1.12%) identify (1.03%) figure (0.88%)	0.136	9.99%
	consumption (54.89%) intertemporal (3.64%) consume (3.06%) sumption (2.08%) consump (1.98%) con (1.82%) tion (0.8%) marginal (0.7%)	0.122	9.31%
	steady (18.9%) state (4.97%) transition (4.8%) economy (4.06%) initial (1.83%) dynamic (1.74%) path (1.6%) level (1.34%)	0.114	6.79%
	household (60.75%) hold (2.32%) income (2.01%) house (1.99%) survey (1.29%) total (0.8%) consumption (0.67%) include (0.52%)	0.093	6.65%
	worker (54.81%) wage (2.85%) work (1.62%) high (1.55%) job (1.11%) employ (1.01%) low (0.83%) labor (0.73%)	0.082	7.06%
Econometric Methods	estimate (17.06%) parameter (6.73%) estimation (3.53%) datum (2.45%) specification (2.36%) likelihood (2.09%) function (2.04%) error (1.58%)	0.176	28.44%
	sample (13.09%) estimate (9.6%) datum (4.65%) table (3.67%) report (2.41%) difference (2.1%) mean (2%) observation (1.3%)	0.127	13.03%
	simulation (9.04%) parameter (8.91%) simulate (4.7%) table (3.06%) mean (2.21%) draw (1.96%) method (1.82%) compute (1.76%)	0.112	8.05%
	variable (9.71%) coefficient (6.31%) regression (3.75%) significant (2.25%) table (2.02%) datum (1.64%) estimate (1.15%) study (1.14%)	0.079	21.59%
	distribution (28.53%) inequality (11.51%) mean (3.22%) gini (1.58%) measure (1.51%) equality (1.11%) distributional (0.93%) equal (0.84%)	0.072	7.37%
	equation (54.73%) variable (8.14%) exogenous (1.92%) determine (1.38%) endogenous (1.32%) form (1.31%) equa (0.95%) substitute (0.92%)	0.068	8.42%
Theory	optimal (26.27%) problem (6.15%) function (2.13%) maximize (2.12%) solution (2.08%) objective (1.4%) solve (1.37%) optimality (1.29%)	0.254	8.97%
	equilibrium (43.53%) exist (1.87%) existence (1.09%) unique (0.97%) equilib (0.94%) rium (0.87%) equi (0.82%) librium (0.68%)	0.237	8.05%
	proposition (8.83%) proof (2.72%) equilibrium (2.44%) imply (2.27%) high (2.06%) low (1.56%) decrease (1.5%) hold (1.48%)	0.190	18.76%
	steady (18.9%) state (4.97%) transition (4.8%) economy (4.06%) initial (1.83%) dynamic (1.74%) path (1.6%) level (1.34%)	0.123	6.79%
	effect (19.24%) reduce (2.83%) rise (2.79%) change (2.27%) fall (2.25%) raise (1.78%) reduction (1.52%) affect (1.43%)	0.108	24.86%
	tax (52.09%) taxis (8.95%) rate (5.2%) revenue (3.35%) taxation (2.74%) high (1.34%) base (1.33%) burden (0.95%)	0.106	7.54%

Notes. Within each citing field of economics research, topics are sorted in decreasing order according to the estimated correlations. Only topics with a presence of at least 5% in all applied theory papers are listed. We only consider citations received within the first seven years of publication.

Table 11. Topics cited the most in econometric methods papers by articles from different fields of economics research

Citing field of economics research	Topic	Correlation with citations shares	Share of articles including the topic
Applied	sample (13.09%) estimate (9.6%) datum (4.65%) table (3.67%) report (2.41%) difference (2.1%) mean (2%) observation (1.3%)	0.197	22.31%
	estimate (17.06%) parameter (6.73%) estimation (3.53%) datum (2.45%) specification (2.36%) likelihood (2.09%) function (2.04%) error (1.58%)	0.151	30.10%
	series (16.29%) trend (13.35%) datum (5.59%) time (3.66%) estimate (2.42%) postwar (1.5%) table (1.27%) period (1.17%)	0.125	7.88%
	hypothesis (32.25%) test (11.01%) reject (8.36%) null (7.68%) datum (3.28%) significance (2.48%) alternative (1.98%) evidence (1.79%)	0.099	12.04%
	cointegrate (3.91%) phillips (3.74%) cointegration (3.68%) regression (2.29%) vector (2.01%) series (1.94%) long (1.51%) granger (1.51%)	0.085	8.22%
	price (69.94%) change (1.36%) average (0.52%) low (0.48%) high (0.38%) different (0.33%) cost (0.33%) demand (0.28%)	0.078	5.11%
Applied Theory	estimate (17.06%) parameter (6.73%) estimation (3.53%) datum (2.45%) specification (2.36%) likelihood (2.09%) function (2.04%) error (1.58%)	0.122	30.10%
	probability (44.16%) conditional (3.55%) distribution (2.08%) transition (2.06%) random (1.7%) number (1.65%) expect (1.54%) prob (1.45%)	0.084	7.74%
	price (69.94%) change (1.36%) average (0.52%) low (0.48%) high (0.38%) different (0.33%) cost (0.33%) demand (0.28%)	0.082	5.11%
	simulation (9.04%) parameter (8.91%) simulate (4.7%) table (3.06%) mean (2.21%) draw (1.96%) method (1.82%) compute (1.76%)	0.066	20.69%
	process (14.95%) conditional (5.39%) time (2.54%) series (1.73%) martingale (1.46%) mean (1.21%) journal (1.03%) stochastic (1%)	0.057	20.64%
	restriction (47.42%) restrict (12.32%) impose (8.5%) unrestricted (5.53%) deregulation (3.06%) allow (1.5%) restric (1.47%) restrictive (1.37%)	0.056	5.45%
Econometric Methods	estimator (9.44%) asymptotic (6.43%) assumption (3.12%) distribution (2.78%) sample (1.87%) asymptotically (1.81%) moment (1.34%) statistic (1.2%)	0.351	50.55%
	nonparametric (5.3%) estimation (5.15%) kernel (4.82%) function (4.3%) moment (4.06%) parametric (3.6%) semiparametric (2.81%) estimator (2.34%)	0.263	16.24%
	bound (4.34%) space (3.83%) assumption (3.28%) continuous (2.86%) compact (2.05%) convex (1.83%) function (1.82%) let (1.59%)	0.169	18.39%
	interval (6.72%) confidence (4.68%) bootstrap (3.23%) inference (1.63%) method (1.43%) sample (1.11%) estimate (1.08%) distribution (1.08%)	0.166	27.76%
	theorem (9.85%) proof (5.57%) let (5.29%) lemma (4.77%) assumption (2.6%) define (2.33%) exist (2.22%) imply (2.13%)	0.165	43.53%
	density (31.8%) distribution (13.58%) likelihood (6.38%) maximum (3.8%) censor (3.07%) observation (2.22%) normal (1.59%) function (1.56%)	0.145	12.42%
Theory	solution (9.48%) problem (5.12%) function (2.83%) convex (2.12%) concave (1.96%) point (1.82%) property (1.74%) unique (1.71%)	0.310	16.96%
	optimal (26.27%) problem (6.15%) function (2.13%) maximize (2.12%) solution (2.08%) objective (1.4%) solve (1.37%) optimality (1.29%)	0.251	6.78%
	price (69.94%) change (1.36%) average (0.52%) low (0.48%) high (0.38%) different (0.33%) cost (0.33%) demand (0.28%)	0.112	5.11%
	bound (4.34%) space (3.83%) assumption (3.28%) continuous (2.86%) compact (2.05%) convex (1.83%) function (1.82%) let (1.59%)	0.111	18.39%
	condition (48.16%) sufficient (4.68%) necessary (3.95%) must (2.15%) satisfy (1.91%) initial (1.61%) positive (1.21%) zero (1%)	0.111	10.08%
	function (45.52%) constant (2.39%) form (1.75%) respect (1.25%) derivative (1.21%) homogeneous (1.15%) derive (1.06%) partial (1.05%)	0.109	20.07%

Notes. Within each citing field of economics research, topics are sorted in decreasing order according to the estimated correlations. Only topics with a presence of at least 5% in all econometric methods papers are listed. We only consider citations received within the first seven years of publication.

Table 12. Topics cited the most in theory papers by articles from different fields of economics research

Citing field of economics research	Topic	Correlation with citations shares	Share of articles including the topic
Applied	income (67.75%) low (1.1%) high (0.86%) level (0.85%) come (0.84%) total (0.66%) equal (0.6%) disposable (0.53%)	0.088	5.17%
	country (38.49%) world (4.87%) develop (3.23%) international (3.04%) high (0.96%) development (0.84%) try (0.72%) coun (0.71%)	0.078	5.57%
	worker (54.81%) wage (2.85%) work (1.62%) high (1.55%) job (1.11%) employ (1.01%) low (0.83%) labor (0.73%)	0.077	5.59%
	wage (70.29%) high (1.18%) low (0.74%) hourly (0.67%) level (0.61%) worker (0.5%) determination (0.48%) labor (0.47%)	0.076	5.62%
	market (76.39%) structure (0.97%) mar (0.64%) ket (0.52%) effect (0.37%) clearing (0.29%) kets (0.29%) different (0.28%)	0.064	6.90%
	effect (19.24%) reduce (2.83%) rise (2.79%) change (2.27%) fall (2.25%) raise (1.78%) reduction (1.52%) affect (1.43%)	0.063	21.77%
Applied Theory	effect (19.24%) reduce (2.83%) rise (2.79%) change (2.27%) fall (2.25%) raise (1.78%) reduction (1.52%) affect (1.43%)	0.102	21.77%
	proposition (8.83%) proof (2.72%) equilibrium (2.44%) imply (2.27%) high (2.06%) low (1.56%) decrease (1.5%) hold (1.48%)	0.097	34.78%
	steady (18.9%) state (4.97%) transition (4.8%) economy (4.06%) initial (1.83%) dynamic (1.74%) path (1.6%) level (1.34%)	0.071	5.75%
	wage (70.29%) high (1.18%) low (0.74%) hourly (0.67%) level (0.61%) worker (0.5%) determination (0.48%) labor (0.47%)	0.063	5.62%
	policy (68.23%) maker (1.31%) implement (0.77%) affect (0.49%) economy (0.45%) different (0.4%) adopt (0.35%) outcome (0.31%)	0.057	5.50%
	worker (54.81%) wage (2.85%) work (1.62%) high (1.55%) job (1.11%) employ (1.01%) low (0.83%) labor (0.73%)	0.045	5.59%
Econometric Methods	function (45.52%) constant (2.39%) form (1.75%) respect (1.25%) derivative (1.21%) homogeneous (1.15%) derive (1.06%) partial (1.05%)	0.121	13.25%
	theory (13.48%) empirical (1.83%) approach (1.55%) theoretical (1.44%) view (1.21%) work (1.07%) concept (1.02%) general (1%)	0.056	15.64%
	bound (4.34%) space (3.83%) assumption (3.28%) continuous (2.86%) compact (2.05%) convex (1.83%) function (1.82%) let (1.59%)	0.051	14.30%
	theorem (9.85%) proof (5.57%) let (5.29%) lemma (4.77%) assumption (2.6%) define (2.33%) exist (2.22%) imply (2.13%)	0.049	33.33%
	income (67.75%) low (1.1%) high (0.86%) level (0.85%) come (0.84%) total (0.66%) equal (0.6%) disposable (0.53%)	0.047	5.17%
	risk (46.22%) aversion (8.21%) averse (2.17%) certainty (1.43%) risky (1.34%) absolute (1.17%) mean (0.99%) decrease (0.93%)	0.044	6.45%
Theory	game (17.82%) player (15.3%) payoff (5.12%) equilibrium (4.65%) play (4.58%) strategy (1.63%) outcome (1.23%) subgame (0.93%)	0.162	9.63%
	equilibrium (43.53%) exist (1.87%) existence (1.09%) unique (0.97%) equilib (0.94%) rium (0.87%) equi (0.82%) librium (0.68%)	0.159	27.05%
	agent (75.54%) example (0.68%) environment (0.44%) receive (0.39%) depend (0.38%) allow (0.38%) denote (0.35%) problem (0.33%)	0.127	7.90%
	payoff (13.71%) ante (4.17%) renegotiation (2.88%) outcome (1.72%) party (1.33%) suppose (1.07%) assumption (1.04%) incomplete (1.04%)	0.118	7.41%
	proposition (8.83%) proof (2.72%) equilibrium (2.44%) imply (2.27%) high (2.06%) low (1.56%) decrease (1.5%) hold (1.48%)	0.105	34.78%
	information (43.51%) inform (3.93%) know (2.82%) reveal (1.93%) uninformed (1.91%) observe (1.74%) informational (1.51%) informa (1.2%)	0.100	7.66%

Notes. Within each citing field of economics research, topics are sorted in decreasing order according to the estimated correlations. Only topics with a presence of at least 5% in all theory papers are listed. We only consider citations received within the first seven years of publication.

6. Summary and conclusions

We built a large corpus of 24,273 economics research articles published in well-regarded general research economics journals, encompassing articles released from 1970 up to and including 2016. Each article was categorized into one of four fields of economics research (applied, applied theory, econometric methods, or theory). Detailed data on citations and references for each article were collected and further enriched using state-of-the-art machine learning and natural language processing techniques. Our study focuses on analyzing trends in the publication and citation of economics research articles, with a specific emphasis on delineating patterns of specialization within fields of economics research. To our knowledge, this article represents the first comprehensive attempt to document specialization within economics to such a significant extent.

When examining differences across fields of economics research, our findings reveal significant disparities in the content of articles, as well as in their citations and references. Regarding specialization, our analysis indicates that certain fields demonstrate an increasing trend towards specialization, while others show the opposite. Specifically, we observe a growing specialization trend in the fields of theory and econometric methods. In contrast, applied papers appear to be moving in the opposite direction. Applied theory stands on the middle ground, and no clear conclusion regarding specialization in this field can be drawn. Below we expand on these results.

Theory has exhibited a narrowing focus on specific research topics since the 1990s, suggesting a tendency towards specialization. The topics experiencing the most pronounced rise since the mid-1990s in theory articles are closely tied to formal mathematical proofs and game theory studies. For this field of economics research, stable patterns are observed in both the share of extramural citations and references. There is also a stable pattern when analyzing the share of citations received from other fields of economics research. However, a slight declining trend is observed when analyzing the share of references made by theory articles to other fields of economics research.

Econometric methods research has also exhibited a narrowing focus on specific research topics since the 1990s. The topics that have shown the strongest growth in econometric methods research articles are related to computational statistics, the asymptotic properties of estimators, and the bounds of estimators. Regarding econometric methods articles, both the share of extramural citations and the share of citations received from other fields of economics research have decreased, suggesting specialization as well. These patterns are accompanied by the fact that econometric methods articles are being published less frequently in general research economics journals and are receiving fewer citations compared to applied and applied theory articles (a pattern also observed for theory articles).

On the contrary, applied articles have exhibited an expansion in the diversity of topics covered since the 1990s. The topics experiencing the highest upward trends in this field are closely linked to impact analysis, causal analysis, and experimental economics. Additionally, there has been an increase in the proportion of extramural citations received by applied articles, with a significant rise in citations originating from disciplines such as medicine, psychology, law, and to a lesser extent, education. The topics most strongly correlated with extramural citations are predominantly present in applied papers. Furthermore, citations coming from other fields of economics research have increased during the period analyzed. By 2016, applied was ranked the most or second most cited field by any other field of economics research. These trends align with the increased publication frequency of applied articles in general research economics journals and their higher citation rates compared to theory articles.

The case of applied theory articles is somewhat ambiguous. While this field has shown an expansion in the diversity of topics covered since the 1990s, this has not resulted in an increase in the proportion of extramural citations received by applied theory articles, or in the share of citations originating from other fields of economics research. In fact, applied theory articles receive the least citations from fields of study other than economics. Although citation patterns indicate a potential specialization in this field, the upward trends in the breadth of topics covered suggest the opposite. This makes it challenging to make a definitive statement regarding specialization in this particular field of economics research.

Data Availability Statement

The data and scripts that support the findings of this study are openly available in openICPSR at https://doi.org/10.3886/E198921, reference number 198921 (Galiani, Gálvez, and Nachman 2024).

References

Amami, Maha, Gabriella Pasi, Fabio Stella, and Rim Faiz. 2016. “An LDA-Based Approach to

Scientific Paper Recommendation.” In *Natural Language Processing and Information Systems*, edited by Elisabeth Métais, Farid Meziane, Mohamad Saraee, Vijayan Sugumaran, and Sunil Vadera, 200–210. Cham: Springer International Publishing.

Anauati, Victoria, Sebastian Galiani, and Ramiro H. Gálvez. 2016. “Quantifying the Life Cycle of Scholarly Articles across Fields of Economic Research.” *Economic Inquiry* 54 (2): 1339–55. https://doi.org/10.1111/ecin.12292.

———. 2020. “Differences in Citation Patterns across Journal Tiers: The Case of Economics.” *Economic Inquiry* 58 (3): 1217–32. https://doi.org/10.1111/ecin.12867.

Anderson, Katharine A, and Seth Richards-Shubik. 2022. “Collaborative Production in Science: An Empirical Analysis of Coauthorships in Economics.” *The Review of Economics and Statistics* 104 (6): 1241–55. https://doi.org/10.1162/rest_a_01025.

Andrews, Donald W K, and Patrik Guggenberger. 2014. “A Conditional-Heteroskedasticity-Robust Confidence Interval for the Autoregressive Parameter.” *The Review of Economics and Statistics* 96 (2): 376–81. http://www.jstor.org/stable/43554937.

Angrist, Joshua, Pierre Azoulay, Glenn Ellison, Ryan Hill, and Susan Feng Lu. 2017. “Economic Research Evolves: Fields and Styles.” *American Economic Review* 107 (5): 293–97. https://doi.org/10.1257/aer.p20171117.

———. 2020. “Inside Job or Deep Impact? Extramural Citations and the Influence of Economic Scholarship.” *Journal of Economic Literature* 58 (1): 3–52. https://doi.org/10.1257/jel.20181508.

Backhouse, Roger, and Beatrice Cherrier. 2014. “Becoming Applied: The Transformation of Economics after 1970.” The Center for the History of Political Economy Working Paper Series No. 2014-15. http://dx.doi.org/10.2139/ssrn.2526274.

Balke, Nathan S. 2000. “Credit and Economic Activity: Credit Regimes and Nonlinear Propagation of Shocks.” *The Review of Economics and Statistics* 82 (2): 344–49. http://www.jstor.org/stable/2646828.

Becker, Gary S, and Kevin M Murphy. 1992. “The Division of Labor, Coordination Costs, and Knowledge.” *The Quarterly Journal of Economics* 107 (4): 1137–60. https://doi.org/10.2307/2118383.

Beltagy, Iz, Kyle Lo, and Arman Cohan. 2019. “SciBERT: A Pretrained Language Model for Scientific Text.” In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3615–20. Hong Kong, China: Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1371.

Biddle, Jeff E, and Daniel S Hamermesh. 2017. “Theory and Measurement: Emergence, Consolidation, and Erosion of a Consensus.” *History of Political Economy* 49 (Supplement): 34–57. https://doi.org/10.1215/00182702-4166251.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades.” *Journal of Political Economy* 100 (5): 992–1026. http://www.jstor.org/stable/2138632.

Blei, David M, Andrew Y Ng, and Michael I Jordan. 2003. “Latent Dirichlet Allocation.” *The Journal of Machine Learning Research* 3: 993–1022.

Brueckner, Markus, and Hannes Schwandt. 2015. “Income and Population Growth.” *The Economic Journal* 125 (589): 1653–76. http://www.jstor.org/stable/24737993.

Card, David, and Stefano DellaVigna. 2013. “Nine Facts about Top Journals in Economics.” *Journal* of *Economic Literature* 51 (1): 144–61. https://doi.org/10.1257/jel.51.1.144.

Chiappori, Pierre-André, and Steven D Levitt. 2003. “An Examination of the Influence of Theory and Individual Theorists on Empirical Research in Microeconomics.” *American Economic Review* 93 (2): 151–55. https://doi.org/10.1257/000282803321946967.

Currie, Janet, and Mark Stabile. 2003. “Socioeconomic Status and Child Health: Why Is the Relationship Stronger for Older Children?” *The American Economic Review* 93 (5): 1813–23. http://www.jstor.org/stable/3132154.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. “BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding.” In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–86. Minneapolis, Minnesota: Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423.

Dirac, Paul Adrien Maurice. 1940. “The Relation between Mathematics and Physics.” *Proceedings of the Royal Society of Edinburgh* 59: 122–129. https://doi.org/10.1017/S0370164600012207.

Dueker, Michael, and Richard Startz. 1998. “Maximum-Likelihood Estimation of Fractional Cointegration with an Application to U.S. and Canadian Bond Rates.” *The Review of Economics and Statistics* 80 (3): 420–26. http://www.jstor.org/stable/2646750.

Duhem, Pierre. 1976. “Physical Theory and Experiment.” In , edited by Sandra G Harding, 1–40. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-010-1863-0_1.

Dunham, William. 1994. *The Mathematical Universe: An Alphabetical Journey Through the Great Proofs, Problems, and Personalities*. New York: Wiley.

Evans, George W. 1991. “Pitfalls in Testing for Explosive Bubbles in Asset Prices.” *The American Economic Review* 81 (4): 922–30. http://www.jstor.org/stable/2006651.

Galiani, Sebastian, and Ramiro H. Gálvez. 2019. “An Empirical Approach Based on Quantile Regression for Estimating Citation Ageing.” *Journal of Informetrics* 13 (2): 738–50. https://doi.org/https://doi.org/10.1016/j.joi.2019.03.014.

Galiani, Sebastian, Ramiro H. Gálvez, and Ian Nachman. 2024. “ECIN Replication Package for ‘Specialization Trends in Economics Research: A Large-Scale Study Using Natural Language Processing and Citation Analysis.’” Ann Arbor, MI: Inter-university Consortium for Political and Social Research. https://doi.org/https://doi.org/10.3886/E198921.

Gálvez, Ramiro H. 2017. “Assessing Author Self-Citation as a Mechanism of Relevant Knowledge Diffusion.” *Scientometrics* 111 (3). https://doi.org/10.1007/s11192-017-2330-1.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. “Text as Data.” *Journal of Economic Literature* 57 (3): 535–74. https://doi.org/10.1257/jel.20181020.

Hall, David, Daniel Jurafsky, and Christopher D Manning. 2008. “Studying the History of Ideas Using Topic Models.” In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, 363–71. Honolulu, Hawaii: Association for Computational Linguistics. https://aclanthology.org/D08-1038.

Hamermesh, Daniel S. 2013. “Six Decades of Top Economics Publishing: Who and How?” *Journal of Economic Literature* 51 (1): 162–72. https://doi.org/10.1257/jel.51.1.162.

———. 2018. “Citations in Economics: Measurement, Uses, and Impacts.” *Journal of Economic Literature* 56 (1): 115–56. https://doi.org/10.1257/jel.20161326.

Hamermesh, Daniel S, and Lea-Rachel Kosnik. 2023. “Why Do Older Scholars Slow Down?” Working

Paper Series. https://doi.org/10.3386/w31175.

Harrow, Keith. 1979. “Theoretical and Applied Computer Science: Antagonism or Symbiosis?” *The American Mathematical Monthly* 86 (4): 253–60. http://www.jstor.org/stable/2320741.

Hastie, T, R Tibshirani, and J H Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd ed. Springer Series in Statistics. New York, NY, USA: Springer.

Heckman, James J, and Sidharth Moktan. 2020. “Publishing and Promotion in Economics: The Tyranny of the Top Five.” *Journal of Economic Literature* 58 (2): 419–70. https://doi.org/10.1257/jel.20191574.

Hu, Diane J. 2009. “Latent Dirichlet Allocation for Text, Images, and Music.” *University of California, San Diego*. http://cseweb.ucsd.edu/~dhu/docs/research_exam09.pdf.

Jackson, C Kirabo. 2010. “Do Students Benefit from Attending Better Schools? Evidence from Rule-Based Student Assignments in Trinidad and Tobago.” *The Economic Journal* 120 (549): 1399–1429. http://www.jstor.org/stable/40929755.

Jurafsky, Daniel, and James H Martin. 2009. *Speech and Language Processing (2nd Edition)*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc.

Kartal, Melis. 2015. “A Comparative Welfare Analysis of Electoral Systems with Endogenous Turnout.” *The Economic Journal* 125 (587): 1369–92. http://www.jstor.org/stable/24737576.

Kinney, Rodney, Chloe Anastasiades, Russell Authur, Iz Beltagy, Jonathan Bragg, Alexandra Buraczynski, Isabel Cachola, et al. 2023. “The Semantic Scholar Open Data Platform.” https://doi.org/https://doi.org/10.48550/arXiv.2301.10140.

Koch, Sigmund. 1973. “Theory and Experiment in Psychology.” *Social Research* 40 (4): 691–707. http://www.jstor.org/stable/40970161.

Kosnik, Lea-Rachel. 2015. “What Have Economists Been Doing for the Last 50 Years? A Text Analysis of Published Academic Research from 1960–2010.” *Economics* 9 (1): 20150013. https://doi.org/doi:10.5018/economics-ejournal.ja.2015-13.

———. 2016. “In Tandem or Out of Sync? Academic Economics Research and Public Policy Measures.” *Contemporary Economic Policy* 34 (1): 190–202. https://doi.org/https://doi.org/10.1111/coep.12117.

Krishnan, Armin. 2009. “What Are Academic Disciplines? Some Observations on the Disciplinarity vs. Interdisciplinarity Debate.” NCRM Working Paper Series. National Centre for Research Methods. http://eprints.ncrm.ac.uk/783/1/what_are_academic_disciplines.pdf.

Lee, Jihong. 2008. “Unforeseen Contingency and Renegotiation with Asymmetric Information.” *The Economic Journal* 118 (528): 678–94. http://www.jstor.org/stable/20108817.

Levitt, Steven D. 1998a. “Juvenile Crime and Punishment.” *Journal of Political Economy* 106 (6): 1156–85. https://doi.org/10.1086/250043.

Levitt, Steven D. 1998b. “Why Do Increased Arrest Rates Appear to Reduce Crime: Deterrence, Incapacitation, or Measurement Error?” *Economic Inquiry* 36 (3): 353–72. https://doi.org/https://doi.org/10.1111/j.1465-7295.1998.tb01720.x.

Long, James E, David W Rasmussen, and Charles T Haworth. 1977. “Income Inequality and City Size.” *The Review of Economics and Statistics* 59 (2): 244–46. https://doi.org/10.2307/1928824.

Longo, Giuseppe, and Ana M Soto. 2016. “Why Do We Need Theories?” *Progress in Biophysics and Molecular Biology* 122 (1): 4–10. https://doi.org/https://doi.org/10.1016/j.pbiomolbio.2016.06.005.

Mailath, George J, Larry Samuelson, and Jeroen M Swinkels. 1993. “Extensive Form Reasoning in Normal Form Games.” *Econometrica* 61 (2): 273–302. https://doi.org/10.2307/2951552.

Martín-Martín, Alberto, Mike Thelwall, Enrique Orduna-Malea, and Emilio Delgado López-Cózar. 2021. “Google Scholar, Microsoft Academic, Scopus, Dimensions, Web of Science, and OpenCitations’ COCI: A Multidisciplinary Comparison of Coverage via Citations.” *Scientometrics* 126 (1): 871–906. https://doi.org/10.1007/s11192-020-03690-4.

Matthews, David. 2021. “Drowning in the Literature? These Smart Software Tools Can Help.” *Nature* 597 (7874): 141–42.

May, Ann Mari, Mary G McGarvey, Yana van der Meulen Rodgers, and Mark Killingsworth. 2021. “Critiques, Ethics, Prestige and Status: A Survey of Editors in Economics.” *Eastern Economic Journal* 47 (2): 295–318. https://doi.org/10.1057/s41302-021-00188-6.

Mustard, David B. 2003. “Reexamining Criminal Behavior: The Importance of Omitted Variable Bias.” *The Review of Economics and Statistics* 85 (1): 205–11. https://doi.org/10.1162/rest.2003.85.1.205.

Neff, Bryan D., and Julian D. Olden. 2010. “Not So Fast: Inflation in Impact Factors Contributes to Apparent Improvements in Journal Quality.” *BioScience* 60 (6): 455–59. https://doi.org/10.1525/bio.2010.60.6.9.

Neral, John, and Jack Ochs. 1992. “The Sequential Equilibrium Theory of Reputation Building: A Further Test.” *Econometrica* 60 (5): 1151–69. https://doi.org/10.2307/2951542.

Panhans, Matthew T, and John D Singleton. 2017. “The Empirical Economist’s Toolkit: From Models to Methods.” *History of Political Economy* 49 (Supplement): 127–57. https://doi.org/10.1215/00182702-4166299.

Phillips, Colin, Phoebe Gaston, Nick Huang, and Hanna Muller. 2021. “Theories All the Way Down: Remarks on ‘Theoretical’ and ‘Experimental’ Linguistics.” In *The Cambridge Handbook of Experimental Syntax*, edited by GrantEditor Goodall, 587–616. Cambridge Handbooks in Language and Linguistics. Cambridge University Press. https://doi.org/10.1017/9781108569620.023.

Pieters, Rik, and Hans Baumgartner. 2002. “Who Talks to Whom? Intra- and Interdisciplinary Communication of Economics Journals.” *Journal of Economic Literature* 40 (2): 483–509. https://doi.org/10.1257/002205102320161348.

Poincaré, Henri. 1902. “Relations Between Experimental Physics and Mathematical Physics.” *The Monist* 12 (4): 516–43. http://www.jstor.org/stable/27899347.

Putnam, Hilary. 1979. “The ‘Corroboration’ of Theories.” In *Mathematics, Matter and Method: Philosophical Papers*, 2nd ed., 1:250–269. Cambridge University Press. https://doi.org/10.1017/CBO9780511625268.018.

Ramaley, Francis. 1930. “Specialization in Science.” *Science* 72 (1866): 325–26. https://doi.org/10.1126/science.72.1866.325.

Reeves, Scott, Mathieu Albert, Ayelet Kuper, and Brian David Hodges. 2008. “Why Use Theories in Qualitative Research?” *BMJ* 337. https://doi.org/10.1136/bmj.a949.

Řehůřek, Radim, and Petr Sojka. 2010. “Software Framework for Topic Modelling with Large Corpora.” In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 45–50. Valletta, Malta: ELRA.

Roskamp, Karl W. 1977. “Labor Productivity and the Elasticity of Factor Substitution in West German Industries 1950-1960.” *The Review of Economics and Statistics* 59 (3): 366–71.

https://doi.org/10.2307/1925056.

Walsh, John P, and Nancy G Maloney. 2007. “Collaboration Structure, Communication Media, and Problems in Scientific Work Teams.” *Journal of Computer-Mediated Communication* 12 (2): 712–32. https://doi.org/https://doi.org/10.1111/j.1083-6101.2007.00346.x.

Wen, Quan. 1994. “The ‘Folk Theorem’ for Repeated Games with Complete Information.” *Econometrica* 62 (4): 949–54. http://www.jstor.org/stable/2951740.

———. 2002. “A Folk Theorem for Repeated Sequential Games.” *The Review of Economic Studies* 69 (2): 493–512. https://doi.org/10.1111/1467-937X.00214.

Wray, K Brad. 2005. “Rethinking Scientific Specialization.” *Social Studies of Science* 35 (1): 151–64. http://www.jstor.org/stable/25046633.

Yano, Makoto. 1985. “Competitive Equilibria on Turnpikes in a McKenzie Economy, II: An Asymptotic Turnpike Theorem.” *International Economic Review* 26 (3): 661–69. https://doi.org/10.2307/2526712.

Ziman, J M. 1987. *Knowing Everything about Nothing: Specialization and Change in Research Careers*. Cambridge University Press. https://books.google.com.ar/books?id=ZEaCTDOw-N8C.