Macrodynamics of Economics: A Bibliometric History

François Claveau and Yves Gingras

Academic disciplines have grown enormously since the Second World War and economics is no exception. Among the economics articles indexed in Thomson Reuter's Web of Science, around a thousand were published in 1956, while nearly twenty thousand are annually published today. This increase amounts to an average annual growth rate between 5 and 6 percent, which is comparable to the global annual growth of science over the same period (Larsen and von Ins 2010). To be sure, these numbers give only a rough indication of the growth of economics, because the Web of Science does not index all academic articles, but only a subset appearing in the major academic journals around the world. Yet, these rough numbers are sufficient to indicate the magnitude of the transformations that have

Correspondence may be addressed to François Claveau, University of Sherbrooke, Department of Philosophy and Applied Ethics, 2500 Boulevard de l'Université, Sherbrooke, Quebec J1K 2R1 (e-mail: francois.claveau@usherbrooke.ca); and to Yves Gingras, University of Québec at Montréal, Department of History, Case Postale 8888, succursale centre-ville, Montréal, Québec H3C 3P8 (e-mail: gingras.yves@uqam.ca). For technical support, we thank the team of the Observatoire des sciences et des technologies. For thoughtful comments on previous versions, we are grateful to one anonymous referee, to Beatrice Cherrier (who dropped her anonymity in the refereeing process), to Till Düppe, and to participants at the History of Economics Society annual meetings in 2013 (Vancouver) and 2014 (Montréal). We thank Saad Blakime and Timothy D. Bowmans for programming the web application and Aymeric Schneider for research assistance. François Claveau acknowledges the financial support of the Social Sciences and Humanities Research Council of Canada for the research phase that led to this article, and we thank the History of Economics Society (New Initiatives Program) and the University of Sherbrooke for funding the development of the web application.

History of Political Economy 48:4 DOI 10.1215/00182702-3687259

affected the scale of scientific research in general and in economics in particular over the last half century.

Although traditional types of analysis of the development of economics—textual analysis of published documents and archives, interviews, biographical reconstruction of major figures, genealogy of concepts and methods—remain useful, specific tools are worth exploring to grasp the global structure and morphology of economics as it changed over the second half of the twentieth century. This article introduces a combination of such tools to the history of economics. We combine bibliometrics with dynamic network analysis to identify a changing specialty structure in economics from the late 1950s to 2014. We map how specialties have emerged, grown in relation to each other, and (for some of them) disappeared. Our results nicely combine with more narrative accounts of the recent history of economics (e.g., Backhouse 2002, chap. 14; Backhouse and Cherrier 2014; and Morgan and Qin 2015): they confirm certain claims that can be found in the existing literature and identify intriguing patterns that invite more in-depth, qualitative research (see section 5).

Bibliometric data have been used in economics, but essentially for descriptive statistics with evaluative purposes—for instance, rankings of journals and most-cited authors (Kalaitzidakis, Mamuneas, and Stengos 2003; Kodrzycki and Yu 2006; Ritzberger 2008; Stern 2013; Zimmermann 2013; Card and DellaVigna 2013). Only a few articles using bibliometric data from economics have no explicit evaluative goal, focusing rather on the structure of authorship and the methods used in articles in a small number of so-called 'top' journals (Hamermesh 2013), on identifying the characteristics of the most-cited articles (Kim, Morse, and Zingales 2006), or on describing the evolution of the number of papers in a given research domain (e.g., Silva and Teixeira 2008, Silva and Teixeira 2009).

To our knowledge, this article is the first in the field of the history of economics to combine bibliometric data and dynamic network analysis. In addition, our corpus of around 415,000 documents is vastly larger than the corpora used in most other studies on the history of economics.[1] Our

1. The closest research on the history of economics we are aware of is Kelly and Bruestle 2011 in which the proportion of articles in each general *JEL* subject category is tracked through time. One downside of that approach is that it has to accept at the start the *JEL* categories as an appropriate grid to classify economics. Our approach lets the algorithm generate the grid at each period of time. Another downside is that Kelly and Bruestle's approach cannot be exported to disciplines that do not have institutionalized classificatory schemes such as *JEL* codes. In contrast, our article is inserted in a larger project aimed at mapping specialty structures in various disciplines. We also start the analysis in 1956 while they start in 1970.

article follows in the footsteps of a number of studies that have shown that bibliometrics and network analysis are particularly well suited for large-scale studies of the historical development of scientific disciplines and specialties (Garfield 2009; Cole, Cole, and Dietrich 1978; Boyack, Börner, and Klavans 2009; Gingras 2010a, 2010b, 2010c; Gingras and Schinckus 2011). The great advantage of combining bibliometrics and network analysis is the possibility to algorithmically detect and visualize patterns in the research output of large groups. Patterns that might remain hidden or blurred with standard historical methods suddenly come in full light.

The main results and analysis are in section 5. What is uncovered is a structuring of economics around recognizable specialties with some significant changes over time and up to the end of our period. Figure 4 is the most informative representation of our results in an article format. But there is an even more revealing representation in the form of an interactive web platform, which combines figure 4 with much more information.[2] Before presenting and analyzing the results, the next four sections describe the conceptual framework, the general method adopted, the data, and the specific empirical procedure. We have decided to reduce the technical information to a minimum in the article; the interested reader is referred to our technical appendix, which can be found at savoirs.usherbrooke.ca/handle/11143/8886.

1. Conceptual Framework: Economics as a Network of Specialties

In any mature scientific discipline, the scientific community is in fact composed of relatively autonomous subunits (the specialties) in which new knowledge is made public through publications. These different specialties, which emerged over time with the morphological growth and ensuing division of labor within the discipline, are more or less connected to each other, much as the discipline itself is more or less connected to other disciplines making up the global scientific field. Whereas the discipline usually corresponds to the entry-level basic (BA) training in a university department, the specialty corresponds more to the research level learned at the MA and PhD levels and lead to new knowledge production (Hagstrom 1965; Cole and Cole 1973; Whitley 1976; Whitley 2000; Bourdieu 2004; Wray 2005).

2. The web platform is located at epistemologiepratique.recherche.usherbrooke.ca/digital-history/economics/. We thank the History of Economics Society and the University of Sherbrooke for financial support.

Disciplines are not eternal, but they tend to be much more stable and institutionalized—as university departments for example—than specialties (Abbott 2001). We could have looked at the changing links and boundaries between economics and other disciplines such as sociology or psychology. Instead, our analysis here focuses on the changing internal structure of economics and thus reconstructs the dynamics of emerging and dying specialties within a given discipline (economics).

Some clarifications of our notion of specialties are in order. First, the claim that a discipline is divided into more or less autonomous specialties is compatible with the claim that it has a core set of theories or methods that most, if not all, specialties rely on.[3] Differential equations are all over physics, and basic equations, such as Maxwell's equations, are also widely used. Similarly, the biomedical sciences make massive use of randomized control trials, although research in oncology and in psychiatry are most probably not in the same specialty. In economics, a large fraction of research relies on a common set of general econometric methods and rational choice models. One might even say that the sharing of these two elements makes economics hold together as a discipline. Yet, research on the budget allocation of poor households in Bangladesh is quite different from theoretical research on optimal monetary policy.

To pursue this point further, we must note that the idea of a specialty structure is also compatible with the widespread view of economics that it has a core that is then applied to different objects (e.g., Ruggles 1970). At a given point in time, some of the specialties might be properly understood as working on developing the future core. For instance, one might think that, early after the Second World War, research on general equilibrium and in econometric theory produced the core theories and methods that have been applied in the following decades to a vast number of topics—this is, in fact, not far from what we find (see section 5.1). In terms of specialties, this view of economics implies that we should find two specialties working on general equilibrium and econometric theory early after the Second World War. These "core" specialties would first develop in parallel with other specialties, but—for these specialties to become the "core"—a period of diffusion would need to follow. The peripheral specialties would come to rely on this core to various degrees. They might even lose their distinct identity and merge with others in the process. During this process of diffusion, other specialties might emerge that work on

3. We thank Beatrice Cherrier for encouraging us to be more precise on this point.

fundamental topics and strive to influence peripheral specialties, thus becoming the new core (if they succeed). Finally, a core that remains stable for some time can undergo what Robert K. Merton (1968, 35) famously labeled "obliteration by incorporation": researchers, taking this core for granted, will come to refer little to the foundational work, rather focusing on connecting explicitly to more recent research that adapts this core to the concerns of their specialty.[4]

Our second clarification comes from the distinction between specialties and the fields of researchers. We conceive specialties as *cognitive* divisions, rather than organizational divisions among a population of researchers based on their primary professional identities. The implication is that the specialty structure does not map each researcher to a single specialty. In fact, some of our results relate to how much economists straddle the specialties of their discipline in comparison to physicists (see section 5.2).

Third and finally, we take the research article as the transmittable knowledge unit par excellence and thus understand specialties as constituted by a set of articles having a relatively high cognitive similarity compared to the rest of articles in the discipline. Cognitive similarity can be due to the sharing of various aspects: shared empirical objects (e.g., education or finance), shared methods (e.g., time series or cross-section econometrics), shared theoretical traditions (e.g., neoclassical or Sraffian economics). Which aspects are more important in determining the identity of specialties depends on context. One advantage of the method we use is that we do not have to a priori decide which aspect must predominate. As we describe below, we use the method of bibliographic coupling to measure cognitive similarity.

2. Methods

To move from this conceptual understanding of disciplines and specialties to a procedure making it possible to follow their development over time, we have to define a measure of the cognitive similarity of two units. This measure will be used to decide whether or not the units belong to the same specialty.

An informal procedure to assess similarity is to rely on insider knowledge: scientists from the discipline often use their trained judgment to evaluate how closely related various bits of research are. Although this

4. In future work, we plan to identify, through the relative dispersion of citations, the specialties that can most properly be interpreted as "core."

informal procedure is often at the source of standard classifications of disciplines and fields, it comes with likely biases (see Cherrier 2015 for a history of the difficult task of updating *JEL* codes). First, it judges similarity from the available categories for distinguishing specialties. It risks being backward-looking for recent research and might thus force the newest publications into categories that do not correspond anymore to the state of the field. Even if it tries to be forward-looking, determining what is likely to become important is a difficult prognosis. For older research, it risks being anachronistic: without sound historical knowledge, one might try to impose current categories that did not apply then. Second, it is vulnerable to strategic claims of similarity and distinctness. Scientists have personal interests at stake in the identification of research that is purportedly close or distant to their own. It might be valuable to be associated to some favorably perceived research or to be perceived as unrelated to other research. It might be even more important for a group to be recognized as distinct and have, for instance, its own *JEL* code. This boundary work (Gieryn 1983) inside science contaminates the informal assessment of similarity.

At the world level of the discipline where thousands of units are involved, approaches fully relying on insider knowledge of the discipline are even less appealing. Our alternative procedure starts by defining the knowledge units of a *discipline* as the documents that are published in journals recognized in that discipline (Meadows 1974; Hagstrom 1965; Merton 1973). To delimit *specialties*, however, no subjective understanding is relied on, since cognitive similarity can be measured from the references contained in the documents of the corpus. Indeed, the method of bibliographic coupling offers a manner to connect two publications on the basis of the similarity of the lists of documents they refer to (Kessler 1963).

Bibliographic coupling is thus used here to measure the similarity between two documents: two documents are considered close to each other when there is a high proportion of overlap in their references. For example, two documents in economics might, at one extreme, contain highly similar work—for instance, the same theoretical literature, the same object, the same method—and thus have almost identical lists of references. At the other extreme, two documents can both be published in economics journals—thus being in the same discipline—but have little else in common. These two documents might have not a single shared reference. They are then probably associated with two very distant specialties of the discipline. The exact measure of similarity between two documents that we use is the number of shared references divided by the square root of the prod-

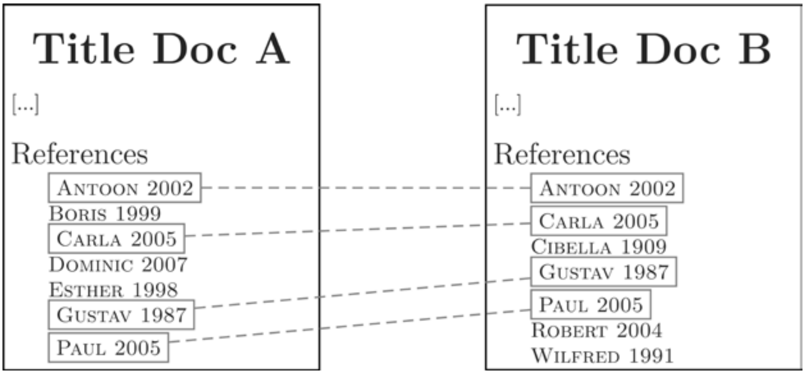

Figure 1 Two documents, each having seven references, with their four shared references highlighted. The cosine measure is calculated as 4 divided by the square root of (7 x 7), or 4/7.

uct of the number of references in each document—that is, the widely used cosine measure (Salton and McGill 1983). Figure 1 illustrates the bibliographic coupling of two documents.[5]

A fundamental assumption of bibliometrics is that bibliographic coupling is an appropriate empirical indicator of the cognitive similarity of two documents. The culture and functions of citations in science offer a solid justification for this assumption. Some key functions of citations are to recognize intellectual priority (Kaplan 1965) and persuade the targeted audience of the novelty and validity of one's research (Gilbert 1977). To perform these functions, authors selectively pick their references based on what they judge relevant to the cognitive content of their paper and what they think their audience will find relevant. It is thus likely that two documents belonging to the same specialty will have a higher proportion of references in common than two documents that are more cognitively distant. In sum, the assumption that the similarity in references is an appropriate indicator of proximity and of belonging to a given specialty is well supported by what we know about the culture and functions of citations in science.

5. It is worth noting that one could also construct the cognitive structure of the discipline by using a complementary manner of creating connections between documents, namely, the co-citations method. From that point of view, which usually gives results similar to bibliographic coupling, two documents are close to each other when they are frequently cited together in other documents (Small and Griffith 1974; Sullivan, White, and Barboni 1977).

Now that we have an operational definition of cognitive similarity between two documents, we could use different methods to identify the specialties inside the discipline. Multidimensional scaling in association with hierarchical clustering is often used (e.g., Kreuzman 2001), but it is much less effective when the corpus includes thousands of documents. Network analysis is a much more powerful method when used in conjunction with community detection techniques (Wallace, Gingras, and Duhon 2009). The nodes of the networks are documents and the intensity of the connections (the edges) between them is given by the bibliographic coupling of each pair of documents.

We can group nodes (documents) in communities[6] by partitioning the set of all nodes in the network. An appropriate partition should group together nodes that are strongly linked—documents with highly overlapping references—and put in different communities nodes that are not linked—documents with no shared reference. The standard measure for the goodness of a partition is called modularity, and community detection proceeds by attempting to select the partition maximizing this measure (Newman and Girvan 2004). Modularity is further discussed in section 4 and in the technical appendix.

These tools will help us analyze the evolving specialty structure of economics. Each specialty will correspond, at a given time, to a subset of the larger network defined by the discipline as a whole. At a larger scale, the same method could be used to partition the whole field of science in different disciplines (Börner 2010). Instead of doing that, we first define economics using a set of journals already recognized within that discipline and included in the Web of Science database. Then, the different specialties emerging over time are determined in an automated fashion by the similarity of the references contained in the documents. No predetermined classification schemes are used to identify specialties. After the specialty structure is produced, each specialty is named by retrieving the terms in the titles of its documents that are most characteristic of it.

3. Data

The data used in this analysis come from the Web of Science of Thomson Reuters. This database samples the entire academic literature; a subset of

6. "Communities" is the technical name in this literature for the groups of nodes found by the detection algorithm. It is not necessarily identical to a specialty: the latter is made of communities, but some communities might not be specialties. We will define the relevant terms in section 4.

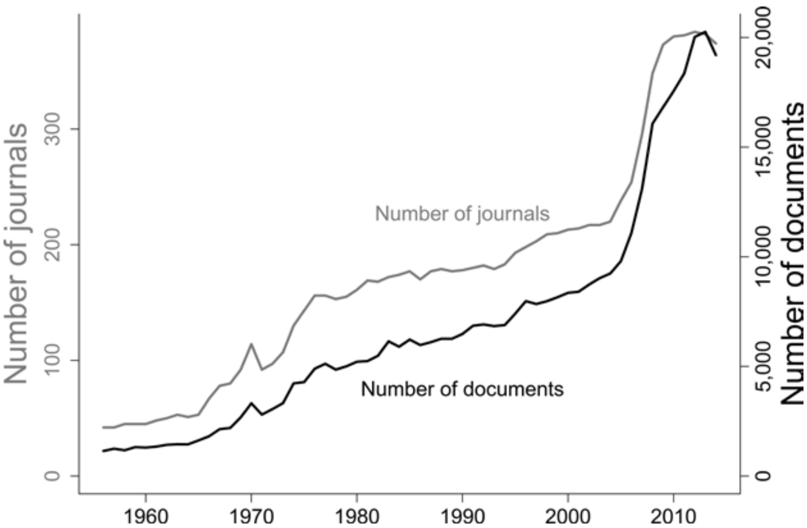

Figure 2 Journals and documents in the corpus. A reduction in quantities is visible at the end of the period because some of the latest journal issues were still in the process of being indexed.

it corresponding to economics is used here. The corpus has been defined in four steps (see the technical appendix for details).

First, a subset of journals in the Web of Science has been selected and divided into mainly economics and partly economics journals. The earliest economics publications indexed go back to 1899 and our sample runs until 2014. Second, documents in these journals have been retained according to different criteria for the two types of journals. The documents that are kept are only articles, notes, and reviews containing references. Figure 2 illustrates the evolution of the number of journals (both mainly and partly economics) and the number of documents published per year in our corpus since 1956.

Third, a starting date for the corpus has been determined based on considerations of data quality. This study makes extensive use of data on references. Unfortunately, the quality of the available data is low for the earliest documents covering the period 1899–1955. For these early years, crucial fields such as first author, publication name, or year of publication are empty for a great number of references. In these circumstances, identifying the document to which the information refers is difficult. Even

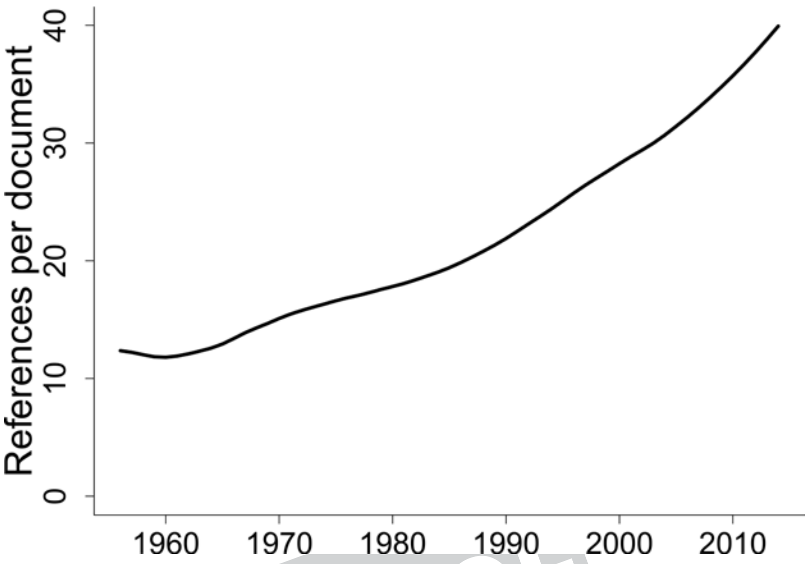

Figure 3 Average number of references per document. The curve is smoothed using local polynomial regression fitting.

after developing a custom identifier for references in order to maximize identification (described in the technical appendix), less than two-thirds of the references are identified prior to 1956. This two-third threshold was chosen as the cutoff point. We will not report results prior to 1956 because the quality of data on references is too low.

In the fourth and last step of our corpus selection, we decided to remove some documents that have vastly more references than the others. This step led to the removal of eight documents. It is a precautionary measure: we do not want our results to be driven by these few peculiar papers. The final corpus is made of 401,278 documents from 536 journals listing a total of 10,381,564 identifiable references.

Our corpus changes in many respects through time. It has already been noted that the number of documents grows. One other crucial variable for us is the average number of references per document, which has also steadily grown (see figure 3). Since the dynamic network that we construct uses documents as nodes and the references to generate the edges, we must be careful in interpreting some of the dynamic properties that we will uncover.

4. Empirical Procedure

Economics changes through time and this should be reflected in the structure of its publications. To follow this evolution, this study defines a window of five years representing a snapshot of publications at a certain moment. This window is moving through time: it starts with the time window 1956–60, then moves to 1957–61, and continues moving forward, one year at a time, until it reaches the last time window, 2010–14. The size of the window (five years) is based on a mixture of considerations. A window of one year is certainly too narrow: given the academic context (e.g., multiple delays in research and in publication), it is typical that papers published a few years back are considered contemporary. Adopting a one-year window also implies that successive snapshots will have no documents in common. In contrast, a multiyear window captures the intuitive idea that a specialty replenishes itself and evolves thanks to new contributions to it, while the older contributions gradually become part of the background. At the other extreme, a window of ten years is probably too wide: a ten-year-old contribution is usually not considered contemporary. In between these extremes, a time window of five years seems reasonable. It acknowledges the delays inherent to research and publishing and it implies that two successive time windows share a little less than 80 percent of their documents. In the technical appendix, we also show the results with a window of three years. Overall, the specialty structure is highly robust to this change in the size of the time window. We will mention in the next section one pattern that is significantly different.

In the dynamic network, the nodes for a specific time window are the documents published in the window and the weighted edges are the number of references that two nodes share normalized with the cosine measure. In this dynamic network, we can apply techniques of automated community detection to group nodes together.

The algorithm used here attempts to maximize "modularity," which is "a measure of the quality of a particular division of a network" (Newman and Girvan 2004, 7). Since *exact* maximization of modularity is computationally hard (Brandes et al. 2008), we use a popular approximate algorithm: the Louvain method (Blondel et al. 2008). This method has been shown to produce partitions with values of modularity at least as high as partitions found with other algorithms. It is also computationally efficient.[7]

7. For an argument in favor of using the Louvain method on *co-citation* data, see Wallace, Gingras, and Duhon 2009.

A great advantage of the Louvain method for this study is that it can easily be adapted to a *dynamic* network.

The Louvain method was originally developed for static networks and thus needs to be adapted for dynamic networks. The key to our modification (suggested by Aynaud and Guillaume 2010, 510–11; Aynaud 2011, sec. 3.5) is to run the algorithm for each time window, but to insert information from the previous time window in the initialization of the algorithm (for details, see the technical appendix). To have prior information even for the first window in our period (i.e., window 1956–60), we start the process as far back as window 1899–1903, but show only results for in-period windows. For each window, the algorithm gives a partition of the documents into communities. A measure is then used to determine whether a community in time window w and one in $w + 1$ are one and the same cluster.[8]

Here is the measure used in this study. Denote by $\mathbf{D}_{\mathbf{w,w+1}}$ the set of documents that are in both time windows, by $\mathbf{d}_i \subset \mathbf{D}_{\mathbf{w,w+1}}$ the subset of these documents that happen to be in community i in window w, and by $\mathbf{d}_{ij} \subset \mathbf{d}_i$ the subset of this subset with the additional property of being in community j in window $w + 1$. The relevant metric is

$$\gamma_{ij} = |\mathbf{d}_{ij}| \, / \, |\mathbf{d}_i|,$$

where $|\mathbf{d}|$ is the cardinality—the number of elements—of $\mathbf{d}$. In short, metric γ_{ij} is giving the fraction of heritable nodes of i that j gets.

In this study, when $\gamma_{ij} \geq 0.65$, community j is considered a child of i.[9] If community j is the child of *only* community i, both communities are taken to be the same cluster persisting through time. We thus have clusters surviving and changing size through time. If a cluster survives through a few windows, we can consider it as an existing *specialty*. But specialties are born and die—the cognitive organization of economics is dynamic. If j is the child of two (or more) communities, j is identified as the realization of a new cluster resulting from the merging of two (or more) clusters. Cluster merging is one type of relation of succession among clusters. In this case,

8. We attribute technical meanings to the terms "community" and "cluster." A community is nothing more than a subset of documents in a given time window, the subsets being determined by the community detection algorithm as applied to the time window. A cluster is a sequence of communities, one per time window for the life length of the cluster. If a cluster lives only one time window, it is identical to a single community.

9. For a justification of this threshold and the other threshold used below, see the technical appendix.

more than one cluster disappears and another cluster is born in the process. This event is particularly significant when some of the merging clusters have been long-lived or when the merged cluster will be long-lived. In these cases, we are observing the reorganization of specialties. When all the involved clusters are short-lived, the situation is better understood as one of weak divisions, not of rapidly changing specialties.

Another type of succession beyond mergers is the splitting of clusters: in this study, if there is no j such that $\gamma_{ij} > 0.65$ and there is at least one j such that $\gamma_{ij} > 0.25$, the cluster of i is taken to have split. Again, the significance of a split regarding the dynamics of specialties hinges on whether it involves long-lived clusters or not.

Finally, we have cases of unilateral births and deaths. If a community is neither a child of one or many previous communities nor the result of a split, we have a case of the birth of a new cluster from scratch. Similarly, if a community i is such that $\gamma_{ij} \leq 0.25$ for all j, this cluster is taken to have simply dispersed, leaving behind neither a child nor a new cluster from its split.

Note that, in order to keep the description focused on important aggregate properties of the discipline, one further restriction is imposed: only the communities regrouping at least 1 percent of the papers in a given time window are considered. The smallest communities are thus left out of the analysis. This restriction helps the analysis avoid being distracted by splits, mergers, and dispersions involving only a tiny and unstable fraction of economics. This restriction combined with the content of our corpus also means that the specialty structure uncovered in this article is a "global" structure dominated by English-American publications, which form the bulk of our corpus. Only regional research strands that are big and well structured have a chance of being detected by our procedure.[10]

The method described so far groups nodes into dynamic clusters by using a modified Louvain method and a metric γ_{ij}. Once we have clusters, we want to have an overview of the documents they contain. What is the content of the research in the various clusters identified? To have such an

10. Following the suggestion of an anonymous referee, we provide in the technical appendix what the main representation looks like if the threshold is a number of documents ($n = 10$) instead of a percentage of the whole discipline. The general specialty structure stays intact, but the representation contains much more noise. We also discuss this alternative restriction below in footnote 15. For the researcher mainly interested in *regional* research strands, this exercise of changing the threshold indicates that a regional corpus and, perhaps, different methods should be used for this endeavor: looking below the 1 percent bar does not give an enlightening picture of what occurred outside the Anglo-American nexus.

overview, we produce a set of keywords for each cluster based on the titles of its documents.

Titles have been preprocessed with standard techniques in text mining: stop words, numbers, and punctuations have been removed and the remaining words have been stemmed. Cluster naming then uses the relative frequencies of stemmed words. More specifically, we use the log-likelihood measure proposed by Paul Rayson and Roger Garside (2000) to retrieve (stemmed) words that are characteristic of each cluster (see the technical appendix for details).

5. Results and Analysis

The most telling representation of our results suitable for a printed article is in figure 4. Each shaded area forming a polygon represents a cluster. The length of the polygon corresponds to how long the cluster survived. As explained above, if the cluster stretches for more than a few time windows, it can be considered a specialty. At every point in time, the height of the polygon corresponds to how large the cluster was in comparison to the whole discipline. The words in each shaded area are the keywords associated with that cluster. The lines connecting polygons indicate that there is a relation of succession among them. If a single cluster is connected to two clusters on its right, this indicates that the first cluster split to give birth to the other two clusters. If two clusters are connected to the same cluster to their right, it indicates that the two clusters merged to give birth to the other cluster. For instance, in row 6, a specialty with keywords "education," "earnings," "labor," and "human" disappeared after 1980–84 to give birth to two short-lived clusters with keywords "wage" and "urban." After only two time windows these clusters merged to give birth to a reunified specialty with keywords "wage," "labor," "employment," and "unemployment." Sometimes (for instance, see row 7), polygons are linearly connected without any branching. This shape happens because the underlying merger or split involves a cluster that is either too small (below 1 percent of the discipline) or too unstable (not surviving two time windows) to appear in the representation.

An even more insightful representation—this time not suitable for a printed article—is an interactive web platform.[11] This platform combines

11. To access it, please visit epistemologiepratique.recherche.usherbrooke.ca/digital-history/economics/.

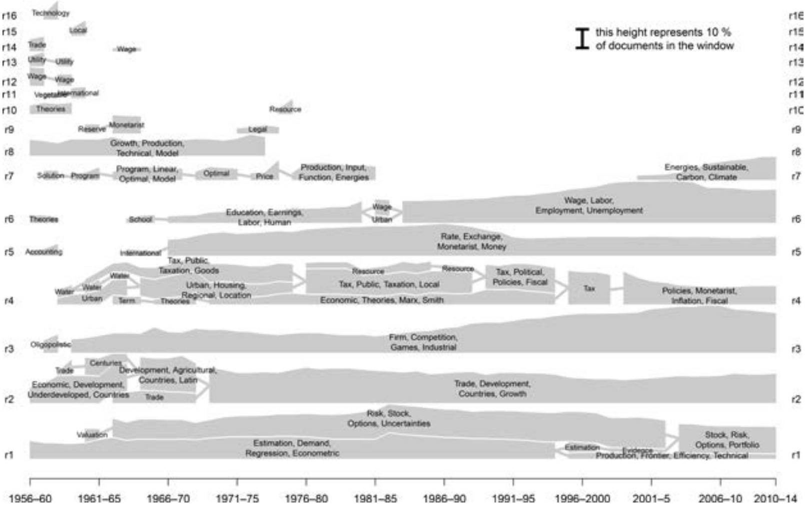

Figure 4 Specialties through time in economics. Only clusters surviving at least two windows are shown. The height of a shaded area represents the relative share of documents per window for the corresponding cluster; the legend at the top right indicates the height corresponding to 10 percent of documents. Connected clusters are either related through mergers or splits. Words for each cluster are the most frequent keywords associated with the cluster during its lifetime. For most clusters surviving at least five windows, the four most frequent keywords are shown; for shorter-lived clusters and a few long-lived but small clusters, only the most frequent keyword is used. Row identifiers are shown on both sides of the image to facilitate referencing. The placement of a cluster on a given row is based only on considerations of readability.

the representation of figure 4, which spans the whole period from 1956 to 2014, with snapshots of the dynamic network in every time window. In figures 5a and 5b, we present the snapshots of the first and last time windows. For each snapshot, nodes represent all the active clusters in the time window. Node sizes are based on relative shares of documents. The words on each node are the keywords associated with that cluster. The thickness of edges—i.e., the lines between nodes—and the proximity of nodes indicate the proximity of clusters in terms of the proportion of shared references. That is, clusters that are close to each other and connected by a

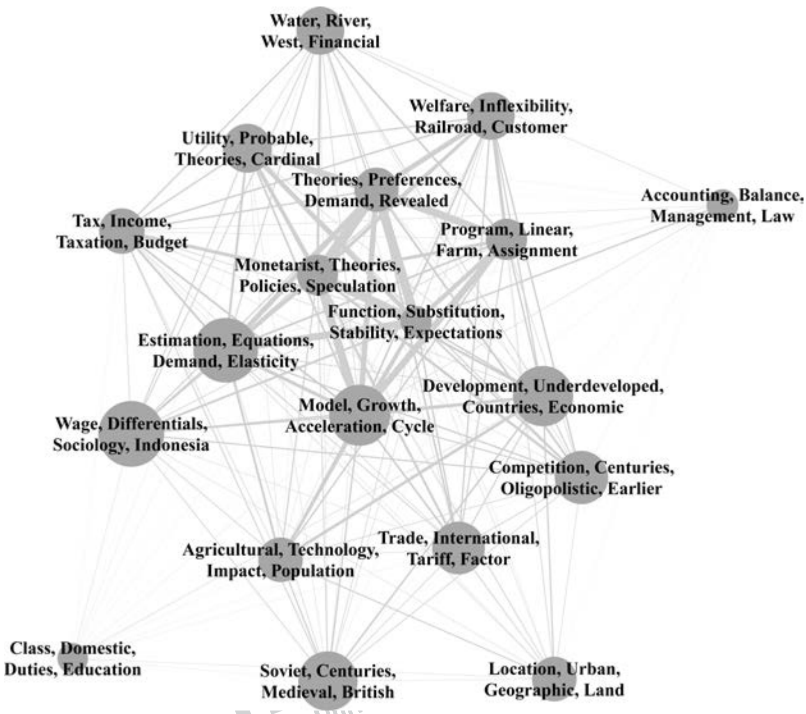

Figure 5a Networks of economics in the first time window. Size of node represents the relative share of documents in the time window; thickness of links and distance between nodes give a rough indication of the proximity of clusters in terms of proportion of shared references. Keywords for each cluster (in each window) are algorithmically extracted from the data. The other snapshots can be visualized and interacted with at epistemologie pratique.recherche.usherbrooke.ca/digital-history/economics/.

thick line are more similar than otherwise. The platform also gives information on each cluster—for instance, life length, characteristic keywords, most commonly cited works.

Since our results are rich, we encourage the reader to explore them with the help of the web platform and the figures and tables in this article. In this section, we want to analyze only a limited number of aspects of the results. We first discuss the ten successions of specialties that are uncovered. We then turn to general properties of the specialty structure. Since the richness of our results might lead one to read *too much* into them, we take care to express what we believe can and cannot be inferred from our results.

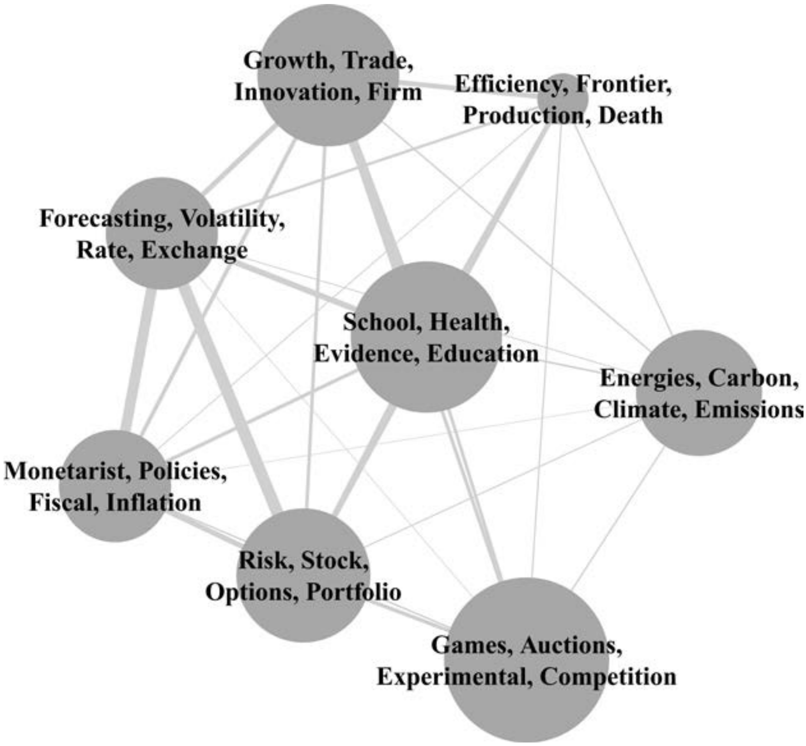

Figure 5b Networks of economics in the last time window. Size of node represents the relative share of documents in the time window; thickness of links and distance between nodes give a rough indication of the proximity of clusters in terms of proportion of shared references. Keywords for each cluster (in each window) are algorithmically extracted from the data. The other snapshots can be visualized and interacted with at epistemologie-pratique.recherche.usherbrooke.ca/digital-history/economics/.

5.1. The Transformation of Specialties

The polygons in our main representation (figure 4) are clusters surviving for at least two time windows. Most of these clusters can be called specialties because they live long enough—we leave it to each reader to decide how long is long enough. Many clusters are connected to each other by relations of merger and split. We call a "succession of specialties" each set of clusters so connected when at least some clusters meet the condition to be specialties. Table 1 identifies the ten successions of specialties of economics from

Table 1 Main Successions of Specialties

	Row in Figure 4	Keywords	First Window	Last Window
1	1	Estimation, Demand, Efficiency, Regression	1956–60	2010–14
2	1	Risk, Stock, Options, Portfolio	1960–64	2010–14
3	2	Development, Trade, Countries, Growth	1956–60	2010–14
4	3	Firm, Competition, Games, Concentration	1957–61	2010–14
5	4	Tax, Public, Economic, Urban	1958–62	2010–14
6	5	Rate, Exchange, Monetarist, Money	1963–67	2010–14
7	6	Labor, Wage, School, Unemployment	1963–67	2010–14
8	7	Program, Linear, Optimal, Production	1957–61	1981–85
9	8	Growth, Production, Technical, Function	1956–60	1973–77
10	7	Energies, Sustainable, Carbon, Climate	2000–04	2010–14

1956 to 2014. Most live for almost the entire period, while the last three span less than half of it. Some are made of few connected clusters, others have a more hectic existence. As will become clear below, even the ones that experience little mergers and splits change significantly throughout the period. We discuss each succession in turn.

1. The first succession, which spans our entire period, can without much risk be labeled *econometric methods*. It stays as a single specialty from time window 1956–60 to window 1994–98. Up to the mid-1980s, it most often refers to econometric textbooks: from Klein 1953 to, most importantly, Theil [1958] 1971 and Johnston 1972. A transition occurs in window 1987–91, when White's 1980 *Econometrica* article on heteroskedasticity jumps to the top of the citation list and stays comfortably there until the end of the specialty. This change is not primarily an evolution from a textbook culture to an article culture since it must be inserted into the larger dynamics of the discipline. In the 1980s, three other specialties become highly reliant on econometric methods (the second succession in row 1 of figure 4, with its longest-lived specialty having keywords "risk," "stock," "options," and "uncertainties," and the successions in rows 5 and 6). The methods developed also become more cognitively specific, for instance with a neater division between time-series econometrics and microeconometrics. As econometrics became "endemic, in various different forms, in the subfields of scientific economic research" (Morgan and Qin 2015, 10), it dispersed itself in our dynamic network. When the specialty splits after window 1994–98, it gives birth to two small communi-

ties. The shortest lived, which continues citing White 1980 abundantly, comes after window 1997–2001 to merge with the succession in financial economics (to be discussed next). The other specialty remains small and focused on statistically measuring efficiency (referencing most often Farrell 1957; Charnes, Cooper, and Rhodes 1978; and Aigner, Lovell, and Schmidt 1977). In sum, at the end of the period, there is not anymore a *general purpose* econometric specialty.

2. The second succession in row 1 of figure 4 consistently has "stock" and "risk" among its keywords. It therefore seems totally appropriate to call it *financial economics*. The succession stays as a single specialty for most of the period. In its earliest time windows, Harry M. Markowitz's *Portfolio Selection* (1959) is the document with the most references. In the 1970s, Gérard Debreu's *Theory of Value* (1959) takes over the first position. In the 1980s and the 1990s, the two most referenced documents become Michael C. Jensen and William H. Meckling's "Theory of the Firm" (1976) and Fischer Black and Myron Scholes's "The Pricing of Options and Corporate Liabilities" (1973). In the 1990s, however, Daniel Kahneman and Amos Tversky's "Prospect Theory" (1979) creeps to the third rank. After a small specialty reorganization in 2003–7 (merger with the remnants of econometrics), this founding article of new behavioral economics becomes the most-cited document of the specialty,[12] closely followed by Eugene F. Fama and Kenneth R. French's "Common Risk Factors in the Returns on Stocks and Bonds" (1993).

3. In row 2 of figure 4, we have a succession of specialties much focused on "development," "trade," and "growth." Up to the 1970s, the specialties on economic development—citing, for instance, Albert O. Hirschman's *Strategy of Economic Development* (1958) and W. W. Rostow's *Stages of Economic Growth* (1960)—are distinct from the specialties on trade, which cite documents such as Harry G. Johnson's *International Trade and Economic Growth* (1958). The two lines of specialties merge in time window 1969–73. Hirshman's book remains highly cited in the 1970s while the specialty is by far the largest of the discipline with an average of 17 percent of the documents per window. In the 1980s, the specialty slowly shrinks in relative size while it shifts to citing most often R. I. McKinnon's *Money and Capital in Economic Development* (1973). In the

12. The evolution of citations to Kahneman and Tversky's article exhibits the "Nobel effect" (Gingras 2010b, 166): after Kahneman received the Bank of Sweden Prize in Economic Sciences in 2002, the annual number of citations to the article kept growing instead of exhibiting the usual radioactive decay shape.

late 1980s, a sudden shift in focus occurs with the rise of the new growth literature. Paul M. Romer's "Increasing Returns and Long-Run Growth" (1986) and Robert E. Lucas's "On the Mechanics of Economic Development" (1988) swiftly install themselves as the most-referenced documents by a wide margin, and Robert J. Barro and Xavier I. Sala-i-Martin's textbook *Economic Growth* (1995) takes the lead in the late 1990s. One noteworthy aspect of this shift to growth is that the previous literature on growth is in a separate succession that died in the mid-1970s (see figure 4, row 8, and our discussion of this specialty below). It is only in the last few windows of the period that significant changes occur in the most-referenced documents. Most importantly, Marc J. Melitz's "Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity" (2003) sits in first position since window 2008–12.

4. In row 3 of figure 4, we find an extremely long-lived specialty with keywords "firm," "competition," "games," and "industrial." Its relative size stays modest from its birth in the late 1950s to the late 1970s, hovering around an average size of 8 percent of the discipline. In its earliest time windows, it refers most often to Joe Bain's *Barriers to New Competition* (1956). Frederic M. Scherer's *Industrial Market Structure and Economic Performance* (1970) becomes more popular in the 1970s. We undoubtedly have here a specialty we can label *industrial economics*. The 1980s is a period of change for this specialty, which starts significantly growing in relative size. Toward the end of the 1980s, new institutional economics has gained predominance, with two publications by Oliver E. Williamson as most-referenced documents (Williamson 1975 and 1985). Noteworthy is also the first appearance in time window 1986–90 of the term "games" in the specialty's keywords. This shift in keywords supports the claim that industrial organization is the specialty where game theory first consolidated itself in economics (Backhouse 2002, 265). But the specialty mutated quite significantly in the process such that it cannot solely be identified with industrial economics in the most recent time windows. This topic undoubtedly remains with, for instance, Jean Tirole's *Theory of Industrial Organization* (1988) being among the three most-referenced documents up to time window 2007–11. Other prominent research avenues include institutional change with Douglass North's *Institutions, Institutional Change and Economic Performance* (1990) being in the top three at the turn of the millennium, and experimental economics. The latter is, in fact, dominant in the last time windows: the specialty's keywords lost most of their industrial-organization flavor (they are "games," "auc-

tions," "experimental," and "competition"), and the three most-referenced documents are in this experimental literature (Fischbacher 2007; Fehr and Schmidt 1999; Bolton and Ockenfels 2000). This diverse and refocused specialty is the largest in the 2000s, with an average of 23 percent of the discipline per time window.

5. In row 4 of figure 4, we discover the most tumultuous succession of specialties. Looking only at the keywords in table 1 (i.e., "tax," "public," "economic," "urban"), one might be tempted to identify this succession with *public finance*, but that is an oversimplification. To facilitate the analysis, we split this succession into shorter sequences of specialties that are listed in table 2. The first sequence—small and short-lived—has a focus on the efficient management of public resources. It merges with the second sequence after time window 1963–67. This second sequence lives up to the late 1970s and, from its most-referenced documents and its keywords, we can say that it centers around the topics of regional and urban economies. It merges with the third sequence, which is what looks most distinctively like a sequence focusing on public finance. Before the merger, this third sequence referred almost three times as much to Richard Musgrave's *Theory of Public Finance* (1959) than to its second most popular reference. The tides change rapidly after the merger. Anthony Downs's *Economic Theory of Democracy* (1957) becomes the most-referenced document in time window 1980–84 and later drops to the second rank after Barro's article "Are Government Bonds Net Wealth?" (1974). In parallel to this transforming specialty, the fourth sequence of specialties, much smaller in size than the third, refers most often to Harold Hotelling's article "The Economics of Exhaustible Resources" (1931) and indeed seems to focus on the topic of exhaustible resources. It merges with the third sequence after time window 1988–92. In the early 1990s, this enlarged third sequence sharpens its focus on what we might call the economic analysis of politics (or, for short, public choice): its keywords now include "political" and it refers most often to Mancur Olson's *Logic of Collective Action* (1965), to Downs 1957, and to Barro 1974.

The story of this succession is not over because a fifth sequence that is born early in the period (window 1962–66) finally interacts in the mid-1990s with the other large specialty. The identity of this fifth sequence is quite unlike what we have seen so far. It starts in the 1960s with a focus on the interest rate structure—referring most often to J. R. Hicks's *Value and Capital* (1939) and to Paul Samuelson's *Foundations of Economic Analysis* (1947). After a split in window 1969–73, which leaves it smaller, the

Table 2 Sequences of Specialties That All Belong to the Family in Row 4 of Figure 4

	Keywords	First Window	Last Window
1	Water, Resource, Public, Cost	1958–62	1963–67
2	Urban, Location, Regional, Housing	1958–62	1975–79
3	Tax, Public, Taxation, Political	1960–64	1994–98
4	Resource, Exhaustible, Extraction, Energies	1976–80	1988–92
5	Economic, Theories, Marx, Smith	1962–66	1994–98
6	Monetarist, Policies, Fiscal, Inflation	1995–99	2010–14

Note: A sequence is defined as a series of linearly connected clusters. A sequence ends when the active cluster has no child, when it merges with a bigger cluster, or, as is the case for the relationship between sequences 3 and 6, when it seems wise to do so given the qualitative changes in the succession of specialties.

sequence refocuses on topics in the *history of economic thought*. Indeed, its most-referenced documents in the early 1970s are Joseph Schumpeter's *History of Economic Analysis* (1954), Mark Blaug's *Economic Theory in Retrospect* (1962), and Adam Smith's *Wealth of Nations* ([1776] 1904). In the mid-1970s, it takes a heterodox—mostly Marxian—flavor: the name "Marx" is part of its keywords from time window 1974–78 to window 1987–91, and Piero Sraffa's *Production of Commodities by Means of Commodities* (1960) is among its most-referenced documents during the same period. This Marxian flavor is replaced in the early 1990s by a Keynesian seasoning: from window 1991–95 to its last window (1994–98), the specialty refers most often to Keynes's *General Theory* (1936) and has "Keynesian" among its keywords. This appearance of Keynes is, from a broader perspective, a displacement: Keynes had been among the most-cited authors up to the mid-1980s in the more mainstream specialty on macro/monetary economics (row 5 of figure 4, to be discussed next).

Finally, the tumultuous succession of specialties in row 4 of figure 4 ends with the merger of the small, historically oriented specialty that we have just discussed with the larger specialty from the third sequence discussed above. This merger gives us the sixth sequence in table 2. In the first time windows of this sequence, we have a clear mix of the previous two sequences, with both Keynes's *General Theory* and Olson's *Logic of Collective Action* among the most-referenced documents. Yet, the sequence develops toward a focus on policy-oriented macroeconomics, with the most-referenced documents becoming Guillermo Calvo's "Staggered Prices in a Utility-Maximizing Framework" (1983), John Taylor's

"Discretion versus Policy Rules in Practice" (1993), and Michael Woodford's *Interest and Prices: Foundations of a Theory of Monetary Policy* (2003). The specialty at the end of the period—with its focus on macroeconomic (chiefly monetary) policy—is quite different from the former specialties in this tumultuous succession. Two elements are worth noting regarding this characteristic. First, it has to be connected with the evolution of the traditional macro/monetary economics succession (figure 4, row 5) toward a widespread use of specific econometric methods (to be discussed next). Second, this succession between a former sequence of specialties focusing on public finance and public choice and a more-recent specialty focused on macroeconomic policy should not be taken too strictly. Indeed, we have here the most important case of a characteristic that is not robust to a change in an assumption of our technique: the time windows of five years rather than three. In the version of our analysis with three-year-long time windows (see the technical appendix), the succession focused on public finance and natural resources ends in the early 1990s and the specialty on macroeconomic policy develops in parallel. The lack of robustness of this characteristic does not mean that it should be *totally* discounted, but the interpretation of this instability will have to wait for a finer analysis of this part of economics than what we can do in this article.

6. The succession of specialties in row 5 of figure 4 is much simpler to characterize than the previous succession. From the mid-1960s to the late 1980s, its main keywords steadily are "monetarist," "rate," "inflation," and "money," and its most-referenced document is Keynes's *General Theory*. The documents just following the *General Theory* in popularity change, however, through this period: up to the mid-1970s, we have Don Patinkin's *Money, Interest, and Prices* (1965), and John Muth's article "Rational Expectations and the Theory of Price Movements" (1961) has its years of glory in the late 1970s before being replaced with Lucas's article "Some International Evidence on Output-Inflation Tradeoffs" (1973) in the early 1980s. We thus have a cognitive specialty clearly dedicated to macroeconomic and monetary issues. Toward the end of the 1980s however, the specialty moves into vector autoregression techniques and, since then, the document receiving by far the most citations is Robert Engle and Clive Granger's article "Co-integration and Error Correction" (1987). This emphasis on an econometric method applied to macroeconomic time series comes together with the development of a parallel specialty on macroeconomic policies (figure 4, row 4; see the previous paragraph).

7. The last succession of specialties that endures for almost the whole period is in row 6 of figure 4. The keywords in table 1 already indicate its identity: "labor," "wage," "school," and "unemployment." It starts small relative to the discipline in the mid-1960s, but grows to become one of the two largest specialties since the mid-1990s. From its birth to window 1976–80, it refers a lot more often to Gary Becker's *Human Capital* (1964) than to anything else. During the same period, its keywords indicate that it focuses on the topics of education and labor markets. In the early 1980s, it integrates the topic of "housing" and diversifies its most-referenced documents with, most notably, Sherwin Rosen's "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition" (1974) and James Heckman's "Sample Selection Bias as a Specification Error" (1979). After the mid-1980s, its keywords attest to a concentration on labor market issues—it is the period when persistently high unemployment in Europe becomes a key policy issue. During the same period, the reorientation toward micro-econometric methods that was already noticeable with the popularity of Heckman 1979 confirms itself. G. S. Maddala's *Limited-Dependent and Qualitative Variables in Econometrics* (1983) is by far the most-referenced document up to window 1996–2000. It then alternatively shares the honor with William Greene's *Econometric Analysis* (1993, first edition) and Heckman 1979. Starting in time window 2003–7, these documents yield to a new generation of econometric texts: Jeffrey M. Wooldridge's *Econometric Analysis of Cross Section and Panel Data* (2002), Manuel Arellano and Stephen Bond's "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations" (1991), and Richard Blundell and Bond's "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models" (1998). The intensity of references to these three documents has been extreme in recent years, as can been seen in table 3: the three documents monopolize the top ten of the most intensively referenced documents in a given year *over our full corpus* (i.e., without restricting to only this specialty). Concurrently to this dominance of new econometric methods for (mainly) panel data, we notice a diversification of the specialty's keywords: "wage" remains, but "school" comes back and "health" appears.[13] Quite significantly, the keyword "evidence" also emerges in the last three time windows. In short, we can confidently assert that this specialty is, toward the

13. In their study of the evolution of the relative frequencies of *JEL* codes, Michael A. Kelly and Stephen Bruestle (2011, table 3) show indeed that the subcategories of "Health" (I10-19) and "Education" (I20-29) have been among the fastest growing of the discipline since the 1990s.

Table 3 Documents Being Referenced Most Intensively in a Given Year

Z-Score	Year Cited	Author	Year	Publication
122.50	2011	Wooldridge, J.	2002	*Econometric Analysis of Cross Section and Panel Data*
122.20	2013	Arellano, M.	1991	*The Review of Economic Studies*
117.60	2012	Wooldridge, J.	2002	*Econometric Analysis of Cross Section and Panel Data*
111.20	2012	Arellano, M.	1991	*The Review of Economic Studies*
109.70	2010	Wooldridge, J.	2002	*Econometric Analysis of Cross Section and Panel Data*
105.50	2014	Wooldridge, J.	2002	*Econometric Analysis of Cross Section and Panel Data*
105.30	2013	Wooldridge, J.	2002	*Econometric Analysis of Cross Section and Panel Data*
102.50	2014	Arellano, M.	1991	*The Review of Economic Studies*
102.20	2009	Wooldridge, J.	2002	*Econometric Analysis of Cross Section and Panel Data*
99.20	2013	Blundell, R.	1998	*Journal of Econometrics*

Note: The first column gives the number of standard deviations from the average number of references to a document in a given year, while the year of the distribution is in the second column. The last three columns give the basic information of the referenced document: first author, year of publication, title of publication (book or journal title).

end of our period, research on a diversity of topics that is united by specific econometric methods.

8. We are left with three successions of specialties that span less than half of our period. The succession on the left of row 7 in figure 4 subsists for a total of twenty-five time windows. Its keywords indicate its focus on *linear programming* and, indeed, it refers most often to *Linear Programming and Economic Analysis* (Dorfman, Samuelson, and Solow 1958) up to the late 1960s. Afterwards, Samuelson's *Foundations of Economic Analysis* (1947) and Hicks's *Value and Capital* (1939, first edition) become its most-referenced documents and its keywords slightly diversify to include "price," "consumption," "welfare," and "production." One might label it "postwar economic theory," as it seems to be work related to the core of economic theory prior to the generalization of game theory. It dies after window 1981–85 and no general, primarily theoretical specialty occurs again. This disappearance and the rising importance of econometric documents in the identity of many specialties as we approach closer to the present support Roger Backhouse and Beatrice Cherrier's (2014, 1) conjecture that "the

discipline became more applied, applied work being accorded a higher status in relation to pure theory than was previously the case."

9. In row 8 of figure 4, we find a specialty active from the earliest time window to the mid-1970s and having keywords such as "growth" and "production." This specialty is clearly the "old" growth theory. As one would expect, it refers most often to Arrow et al.'s article "Capital-Labor Substitution and Economic Efficiency" (1961) and to Solow's "Technical Change and the Aggregate Production Function" (1957). It disappears in the mid-1970s and, as described above, the "new" (endogenous) growth theory appears in the late 1980s in the succession associated with economic development (row 2 of figure 4). This pattern supports an existing narrative that sees the 1970s and early 1980s as a period during which growth was a less prevalent research topic in economics (Boianovsky and Hoover 2009, 3).

10. Finally, a small specialty with no clear succession is born in time window 2000–4 (figure 4, row 7). It gradually grows in relative size to reach 12 percent of the discipline at the end of the period. Its keywords clearly indicate its main topic: "energies," "sustainable," "carbon," and "climate." But the documents to which it refers the most are not on the topic of climate change. Again, they are documents contributing to specific econometric methods: Kenneth Train's *Discrete Choice Methods with Simulation* (2003), Daniel McFadden's "Conditional Logit Analysis of Qualitative Choice Behavior" (1973), and Louviere, Hensher, and Swait's' *Stated Choice Methods: Analysis and Applications* (2000).[14]

5.2. General Properties of the Specialty Structure

So far, we have discussed specialties in turn and not the overall specialty structure of the discipline. When it comes to this structure, we have to be much more careful in our interpretation of the raw results. As we noted at the end of our data section (section 3), many things are changing through time: the number of documents per window, the average number of references per document, the skewness of the distribution of citations, the reliability of the identification of the references. When some aggregate property of the specialty structure shows a clear trend through time, it is often

14. The historical study of climate-change economics is in its infancy. The pattern described in Huet 2015 concords with the one found here.

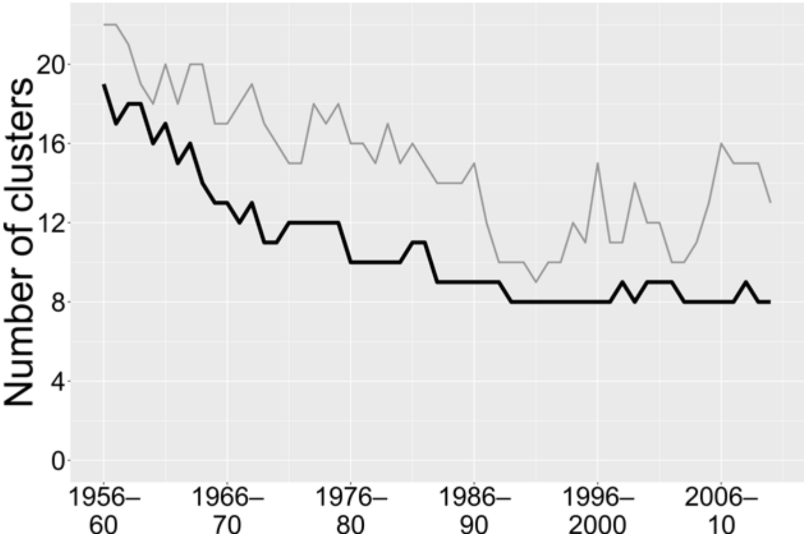

Figure 6 Properties of the specialty structure of economics from window 1956–60 to 2010–14: Number of clusters detected per time window. The two lines differ by the threshold used to exclude the smallest communities from counting as clusters: the black line is the number of clusters when communities below 1 percent of the size of the discipline are excluded; the grey line is when only communities with fewer than ten documents are excluded.

not possible to reliably discard the hypothesis—given the current state of dynamic community detection—that this trend is simply an artifact of some of the changing characteristics of the data just mentioned. Withholding belief seems especially warranted regarding the two aggregate properties that might at first catch the eye: the number and stability of clusters.

Number of clusters. The drop in the number of clusters is immediately evident from the comparison of the two networks in figures 5a and 5b. Figure 5a depicts a network of nineteen clusters in window 1956–60 while this number has dropped to eight in the last time window (figure 5b). Figure 6 presents the number of detected clusters per window over the whole period. For our main specification (black line), the number of clusters slowly declines from the late 1950s to the early 1980s to then stabilize to eight or nine clusters.[15] What should we make of this result? One could conclude that economics research is structured around fewer divisions today. Unfortunately, this conclusion is premature. The main reason is that the technical literature on modularity optimization—a method to which our detection algorithm belongs—finds that, as networks get larger, the method "may favor network partitions with groups of modules combined into larger communities" (Fortunato and Barthélemy 2007, 41). In other words, economics might have as many or perhaps more "divisions" than before, but the method could tend to more frequently lump together smaller groups into bigger specialties as the period unfolds (and the size of the network grows). There are two things to note here. First, it is unclear that this process is actually happening in our network. To ascertain it, we would need independent and reliable access to the actual specialty structure—access we do not have. Second, even if it is happening, it does not invalidate the story we have been telling about the successions of specialties. It simply highlights that there is probably an internal specialty structure to the large specialties that we identify toward the end of the period. In fact, we are quite confident that the two largest specialties since the late 1990s—the first historically linked to industrial economics and the other to labor economics (see figure 4, rows 3 and 6)—have an internal structure that would be worth exploring in future research. As we noted previously in our discussion, these two specialties seem, in recent times, to include diverse bundles of articles that have nevertheless some

15. This result excludes communities that account for less than 1 percent of the discipline. The grey line in figure 6 depicts the number of clusters if the only constraint on size is that a community must count at least ten documents to be registered as a cluster. With this constraint, a similar downward trend in the number of clusters is visible, but more pronounced short-term variations occur since the 1980s. The burst in the number of clusters below 1 percent since window 2005–10 is noteworthy. Are we thus missing something important by restricting ourselves to clusters above 1 percent of the discipline? We think not. The majority of these small clusters since 2005–10 do not even include 50 documents. The largest, with an average of 455 documents per window and surviving six windows, is fully concentrated in one journal—*Actual Problems of Economics*—and its article with the most impact has only five citations. The second largest cluster, with an average of 186 documents over its two window-long life, has more than 80 percent of its documents in two journals—*Custos e Agronegocio* and *Journal of Home Economics Research*—and its highest impact article receives ten citations. In contrast, the smallest cluster above the 1 percent threshold lives sixteen windows, has an average of 1,660 documents per window from window 2005–10 to the end of the period, and contains fifty documents that receive more than one hundred citations (it is the cluster at the bottom right of figure 4 in row 1, having keywords "production," "frontier," "efficiency," and "technical"). The contrast is stark enough to justify using the 1 percent threshold to discriminate between clearly relevant and little relevant clusters in our *macro* history of specialties in economics. We thank an anonymous referee for inciting us to be clearer on the justification for our 1 percent threshold.

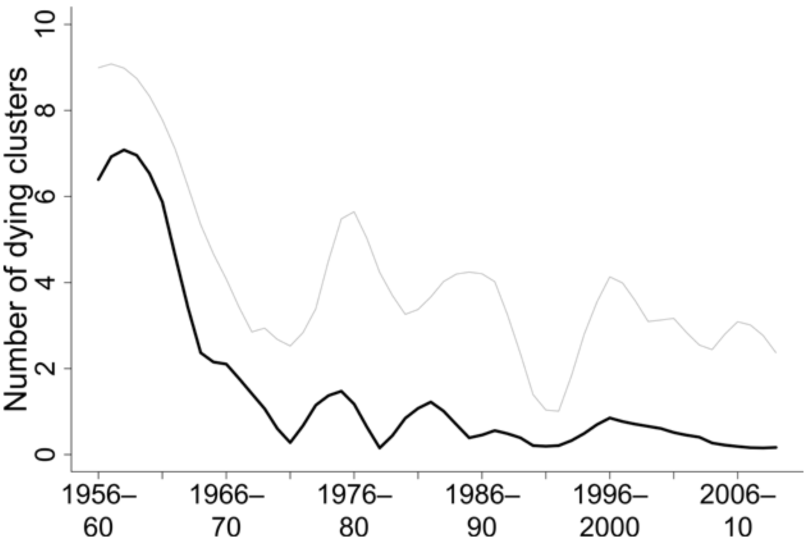

Figure 7 Properties of the specialty structure of economics from window 1956–60 to 2010–14: Instability through time of the specialty structure as measured by the number of clusters that end in each time window. As in figure 6, the two lines stand for different constraints on the size threshold for cluster inclusion (the black line is for our main specification). The curves are smoothed using local polynomial regression fitting.

methods in common. We could test the existence of internal structures by running our community detection algorithm on networks constituted by the documents in only one specialty.

Stability of clusters. The second aggregate property that strikes the eye is the increasing stability of the specialty structure: the left side of figure 4 mostly presents short-lived clusters while the right side mostly exhibits long-lived ones. Figure 7 presents a measure of instability, the number of clusters disappearing after each time window. The smoothed line for our main specification (black line) drops below 2 as early as window 1967–71 to never return. In fact, the structure is so unstable in the earliest time windows that few of the detected clusters should be counted as authentic specialties. We urge the reader not to interpret this result as implying that there were barely no specialties in economics in the late 1950s. That our method fails to detect relatively stable clusters at some point does not imply that there is no specialty. Early in the period, articles had on average fewer references (see figure 3) and these references were poorly indexed in comparison to later ones. Grappling with this limited information, it is plausible that our algorithm fails to find stable patterns, although economics was already cognitively specialized. On this background of limited information, the fact that the method finds some specialties seems more significant than the possibility that it misses some others. Specialties focused on econometrics (figure 6, row 1), on economic development (row 2), and on growth (row 8) have a stable identity as early as the first time window (1956–60). This early stability indicates the extent to which these specialties were cognitively distinct. As the information contained in the references becomes better, our method is able to detect as soon as the mid-1960s all the successions of specialties that will remain until the end of the period—except the one focused on climate change, which clearly did not exist then.

Cluster disconnection. Another aggregate property of the specialty structure that can be tracked through time is the strength of the connections between clusters, as depicted in figure 8. Since this property has no clear trend, we submit that, unlike the two previous properties, it can be interpreted at face value. There have been episodes when clusters where more connected to each other (i.e., lower values of disconnection in figure 8), which means that the boundaries between specialties were thinner. Such episodes occurred at the beginning of the period and then in the mid-1980s. In the mid-1980s, the two specialties having by far the most intense connections between them and to the rest of the network[16] are the econometrics-centered specialty (figure 4, bottom of row 1) and the macro/money specialty (figure 4, row 5). This situation might be interpreted as the first signs of the deep changes of the specialty structure based on distinctions among econometric methods. In the econometrics-centered specialty, the mid-1980s corresponds to the rise of White's 1980 article on heteroskedasticity, which will reach the top three in the citation rankings of all time in economics. In the other central specialty, the mid-1980s corresponds to the shift away from Keynes 1936 as a core reference and toward Engle and Granger's CVAR approach—Engle and Granger 1987 also being in the all-time top three of the discipline. As can be seen from figure 8, the last episode of low disconnection between clusters—in fact, the episode when cluster disconnection reaches its lowest value—coincides with the end of our period, in window 2010–14. At that point, the

16. The strength of the connection between two clusters is simply their weighted edge. The overall centrality of a cluster is given by eigenvector centrality (Bonacich 1987; Newman 2010, 169–72).

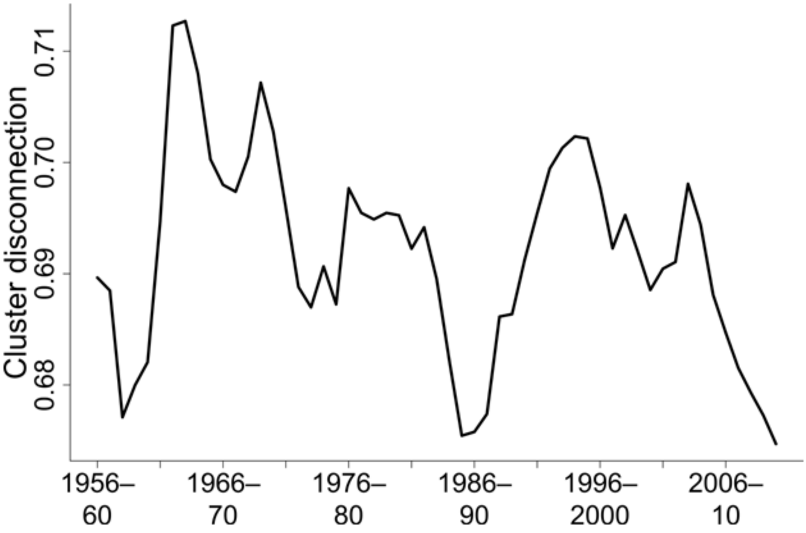

Figure 8 Properties of the specialty structure of economics from window 1956–60 to 2010–14: Degree of cluster disconnection per window. The mathematical definition of cluster disconnection is given in the discussion of modularity in the technical appendix.

specialty that came to focus on CVAR is alone in the category of most central community. Unfortunately, our method is not meant for forecasts. We cannot tell whether this episode is an indication that fundamental restructuring is underway.

Cognitive specialties and individual researchers. The nodes of our dynamic network are documents. The specialties that we detect are cognitive clusters of documents changing with time. In section 1, we distinguished this understanding of specialty from a more organizational understanding where individual researchers would be the nodes. Because a researcher can contribute to more than one cognitive specialty, the structure of the two types of network will not be identical. But how much different is the cognitive structure of economics from its organizational structure? Since our data include all authors, we can assess how much authors straddle cognitive specialties.[17]

17. The results used in the following discussion have not been checked for homonyms—two authors having the same name. Our results might thus be slightly biased, but this bias is most probably not big enough to confound our qualitative results.

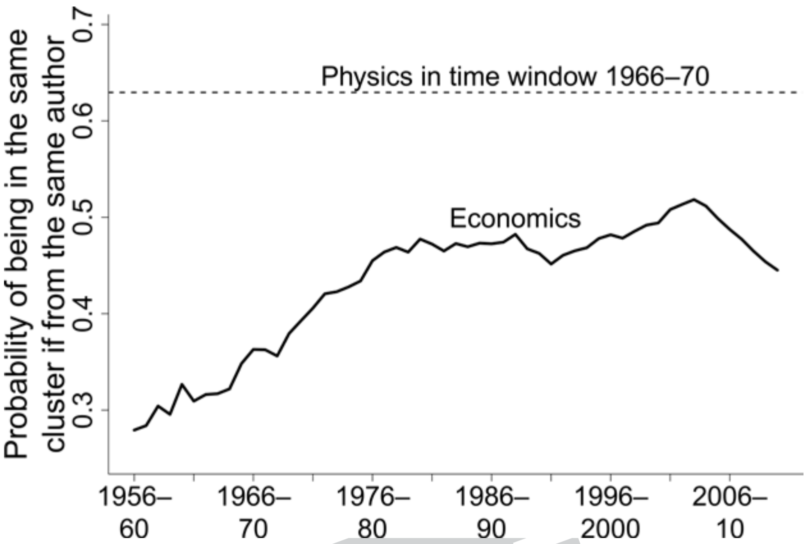

Figure 9 Properties of the specialty structure of economics from window 1956–60 to 2010–14: Probability that two documents with at least one common author are in the same cluster

Perhaps surprisingly, if we take at random from a time window two documents with at least one shared author, the probability that these two documents are in the same cluster surpasses 0.5 in only a few time windows (from 2001–5 to 2004–8). Figure 9 depicts the evolution of this probability. An interesting point of comparison is to take the discipline of physics. We did the same exercise with physics in window 1966–70: the probability that two documents sharing at least one author are in the same cluster is 0.63, a probability much higher than anything found for economics. The comparison is especially relevant because physics in window 1966–70 has a number of authors (89,000) similar to economics toward the end of the period (e.g., 90,500 authors in window 2009–13), but physics is divided into many more clusters (fourteen versus eight or nine in economics). With many more clusters, one would expect that authors would be more likely to straddle clusters. Since our results suggest the opposite, we can confidently conclude that economists are less cognitively specialized in their discipline than physicists of the late 1960s—that is, they tend to contribute to more cognitive specialties in their discipline.

The probability that two documents are in the same cluster given that they share at least one author provides only a partial understanding of the

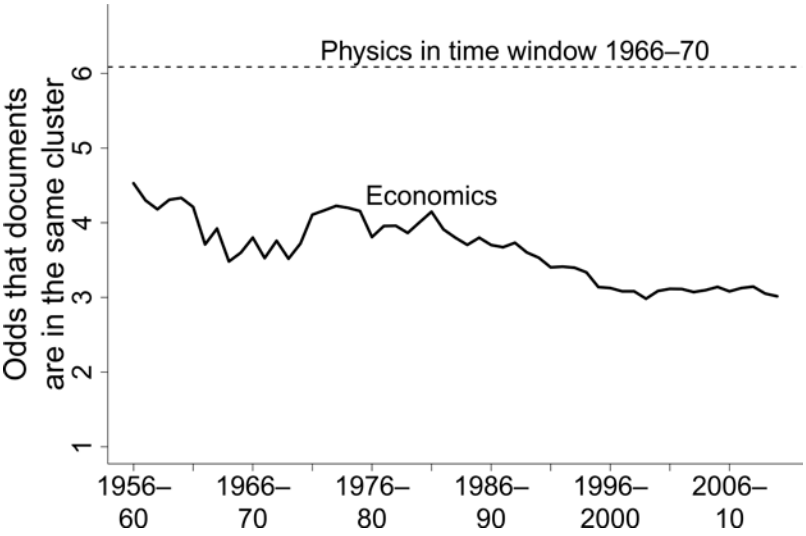

Figure 10 Properties of the specialty structure of economics from window 1956–60 to 2010–14: Ratio of the probability that two documents with at least one common author are in the same cluster (see figure 9) to the same probability given nonoverlapping author lists

relationship between authorship and cognitive specialties. Another question is, How much more likely is it that two documents are in the same cluster given that they share at least one author instead of sharing none? The curve in figure 10 tracks this ratio through time. It is at its highest values in the earliest windows: having at least one shared author back then makes it roughly 4.5 times more likely for two documents to be in the same cluster. Since the number of clusters is extremely high (around eighteen) in these early windows, it is a lot less likely for two documents to be in the same cluster by chance. This component of luck explains the discrepancy between the tendency of these odds to drop with time and the growing probability of being in the same cluster given an author overlap (the numerator in the odds, which is shown in figure 6d). At the end of our period, having an overlap in authors makes it roughly 3 times more likely for two documents to be in the same cluster. The comparison with physics in time window 1966–70 is again instructive: author overlap makes the sharing of a cluster 6.1 times more likely in this discipline.

Finally, instead of looking at how much authors straddle clusters, we can focus on the direct measure of the cognitive similarity between two documents used by our detection algorithm. This is the cosine measure of bibliographic coupling. Figure 11 tracks the ratio of the mean cosine measure between documents sharing at least one author to the mean cosine measure between documents sharing no author at all. A value of 100 would mean that the mean cosine measure for documents with overlapping authors is a hundred times the mean cosine measure for documents with no shared author. For economics over our period, this value steadily grew from 105 to 1,200. This increase through time is mostly attributable to a sharply dropping denominator, since the numerator (the mean cosine measure of documents with overlapping authors) followed a bell-shaped curve. Again as a point of comparison, in physics the mean cosine measure of documents sharing at least one author is 5,675 times higher than the mean cosine measure of documents sharing no author at all, a tall order higher than in economics.

In sum, this analysis of the distinction between the structure of cognitive specialties and an organizational structure of researchers shows that a simple connection between the two structures does not exist. As one would expect, common authorship is one determinant of the similarity of references (figure 11) and, correlatively, of joint membership in a cluster (figure 10). Yet, common authorship in economics could be a stronger predictor of these two properties, as is exemplified by the contrast with physics.

6. Conclusion

This study has provided a map of the macrodynamics of specialties in economics since the late 1950s. After defining a corpus of roughly 415,000 economics documents in the Web of Science, our procedure carved out specialties in an automated fashion by using bibliographic coupling, algorithmic community detection, and text mining. In attempting to understand the global cognitive structure of a discipline as large as economics, this procedure has clear advantages over a method relying mainly on insider knowledge of a discipline.

Our procedure delineates recognizable specialties that exhibit nontrivial dynamics. Grouping specialties together when they are connected by mergers or splits, we arrive at ten successions of specialties (see table 1). Seven of them live for most of the period. The identities of those specialties change quite significantly over time but they originate in well-known fields of economics: econometric methods, financial economics, economic development and international trade, industrial organization, pub-

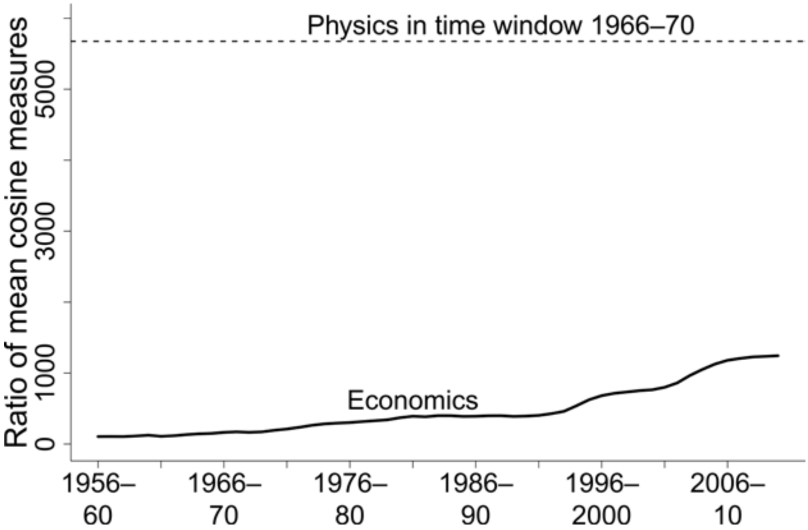

Figure 11 Properties of the specialty structure of economics from window 1956–60 to 2010–14: The ratio of the mean cosine measure between two documents given overlapping author lists to the same mean measure given nonoverlapping author lists

lic finance, macro and monetary economics, and labor economics. Among the three successions of specialties that exist for less than half of the period since the late 1950s, two live in the early time windows—one focuses on what might be called "general, postwar economic theory" and the other focuses on growth (in the footsteps of the Solow growth model). The last specialty emerged in the late 1990s and is concerned with environmental issues—mostly climate change.

Our article is meant as a foray into algorithmic empirical work for *historical* purposes. Much remains to be done in analyzing our results at different scales (in distinguishing countries, for example), in improving the algorithmic methods, and in comparing the specialty dynamics uncovered here with the one in other scientific disciplines. It is time for historians of economics to use the combined resources of large data, algorithmic methods, and computing power. This type of work can complement research using the already established toolbox by providing solid quantitative evidence for existing historical claims and by uncovering intriguing patterns that could then be analyzed in more detail using qualitative approaches.

References

Abbott, Andrew. 2001. *Chaos of Disciplines*. Chicago: University of Chicago Press.

Aigner, Dennis, C. A. Knox Lovell, and Peter Schmidt. 1977. "Formulation and Estimation of Stochastic Frontier Production Function Models." *Journal of Econometrics* 6 (1): 21–37.

Arellano, Manuel, and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies* 58 (2): 277–97.

Arrow, K. J., H. B. Chenery, B. S. Minhas, and R. M. Solow. 1961. "Capital-Labor Substitution and Economic Efficiency." *Review of Economics and Statistics* 43 (3): 225–50.

Aynaud, Thomas. 2011. "Détection de communautés dans les réseaux dynamiques." PhD diss., Université Pierre et Marie Curie.

Aynaud, Thomas, and Jean-Loup Guillaume. 2010. "Static Community Detection Algorithms for Evolving Networks." In *2010 Proceedings of the 8th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, Avignon, France, 508–14.

Backhouse, Roger E. 2002. *The Penguin History of Economics*. London: Penguin Books.

Backhouse, Roger, and Beatrice Cherrier. 2014. "Becoming Applied: The Transformation of Economics after 1970." Center for the History of Political Economy Working Paper 2014-15.http://papers.ssrn.com/abstract=2526274.

Bain, Joe Staten. 1956. *Barriers to New Competition: Their Character and Consequences in Manufacturing Industries*. Cambridge, Mass.: Harvard University Press.

Barro, Robert J. 1974. "Are Government Bonds Net Wealth?" *Journal of Political Economy* 82 (6): 1095–1117.

Barro, Robert J., and Xavier I. Sala-i-Martin. 1995. *Economic Growth*. Boston: McGraw-Hill.

Becker, Gary S. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. Chicago: University of Chicago Press.

Black, Fischer, and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3): 637–54.

Blaug, Mark. 1962. *Economic Theory in Retrospect*. Cambridge: Cambridge University Press.

Blondel, Vincent D., Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. "Fast Unfolding of Communities in Large Networks." *Journal of Statistical Mechanics: Theory and Experiment* 2008 P10008. doi:10.1088/1742-5468/2008/10/P10008.

Blundell, Richard, and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics* 87 (1): 115–43.

Boianovsky, Mauro, and Kevin D. Hoover. 2009. "The Neoclassical Growth Model and Twentieth-Century Economics." In *Robert Solow and the Development of*

Growth Economics, edited by Mauro Boianovsky and Kevin D. Hoover. *History of Political Economy* 41 (supplement): 1–23.

Bolton, Gary E., and Axel Ockenfels. 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review* 90 (1): 166–93.

Bonacich, Phillip. 1987. "Power and Centrality: A Family of Measures." *American Journal of Sociology* 92 (5): 1170–82.

Börner, Katy. 2010. *Atlas of Science: Visualizing What We Know*. Cambridge, Mass.: MIT Press.

Bourdieu, Pierre. 2004. *Science of Science and Reflexivity*. Cambridge: Polity Press.

Boyack, Kevin W., Katy Börner, and Richard Klavans. 2009. "Mapping the Structure and Evolution of Chemistry Research." *Scientometrics* 79 (1): 45–60.

Brandes, U., D. Delling, M. Gaertler, R. Gorke, M. Hoefer, Z. Nikoloski, and D. Wagner. 2008. "On Modularity Clustering." *IEEE Transactions on Knowledge and Data Engineering* 20 (2): 172–88.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Card, David, and Stefano DellaVigna. 2013. "Nine Facts about Top Journals in Economics." *Journal of Economic Literature* 51 (1): 144–61.

Charnes, A., W. W. Cooper, and E. Rhodes. 1978. "Measuring the Efficiency of Decision Making Units." *European Journal of Operational Research* 2 (6): 429–44.

Cherrier, Beatrice. 2015. "Classifying Economics: A History of the JEL Codes." Unpublished paper. CREM, University of Caen.

Cole, Jonathan R., and Stephen Cole. 1973. *Social Stratification in Science*. Chicago: University of Chicago Press.

Cole, Stephen, Jonathan R. Cole, and Lorraine Dietrich. 1978. "Measuring the Cognitive State of Scientific Disciplines." In *Toward a Metric of Science: The Advent of Science Indicators*, edited by Yehuda Elkana, Joshua Lederberg, Robert K. Merton, Arnold Thackray, and Harriet Zuckerman, 209–52. New York: Wiley.

Debreu, G. 1959. *Theory of Value*. New York: Wiley.

Dorfman, Robert, Paul A. Samuelson, and Robert M. Solow. 1958. *Linear Programming and Economic Analysis*. New York: McGraw-Hill.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper.

Engle, Robert F., and C. W. J. Granger. 1987. "Co-integration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55 (2): 251–76.

Fama, Eugene F., and Kenneth R. French. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33 (1): 3–56. doi:10.1016/0304-405X(93)90023-5.

Farrell, M. J. 1957. "The Measurement of Productive Efficiency." *Journal of the Royal Statistical Society: Series A (General)* 120 (3): 253–90.

Fehr, Ernst, and Klaus M Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Quarterly Journal of Economics* 114 (3): 817–68.

Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics* 10 (2): 171–78.

Fortunato, Santo, and Marc Barthélemy. 2007. "Resolution Limit in Community Detection." *Proceedings of the National Academy of Sciences* 104 (1): 36–41.
Garfield, Eugene. 2009. "From the Science of Science to Scientometrics Visualizing the History of Science with HistCite Software." *Journal of Informetrics, Science of Science: Conceptualizations and Models of Science* 3 (3): 173–79.
Gieryn, Thomas F. 1983. "Boundary-Work and the Demarcation of Science from Non-science: Strains and Interests in Professional Ideologies of Scientists." *American Sociological Review* 48 (6): 781–95.
Gilbert, G. Nigel. 1977. "Referencing as Persuasion." *Social Studies of Science* 7 (1): 113–22.
Gingras, Yves. 2010a. "Mapping the Structure of the Intellectual Field Using Citation and Co-citation Analysis of Correspondences." *History of European Ideas* 36 (3): 330–39.
———. 2010b. "Revisiting the 'Quiet Debut' of the Double Helix: A Bibliometric and Methodological Note on the 'Impact' of Scientific Publications." *Journal of the History of Biology* 43 (1): 159–81.
———. 2010c. "The Transformation of Physics from 1900 to 1945." *Physics in Perspective* 12 (3): 248–65.
Gingras, Yves, and Christophe Schinckus. 2011. "The Institutionalization of Econophysics in the Shadow of Physics." *Journal of the History of Economic Thought* (forthcoming).
Greene, William H. 1993. *Econometric Analysis*. Upper Saddle River, N.J.: Prentice-Hall.
Hagstrom, Warren O. 1965. *The Scientific Community*. New York: Basic Books.
Hamermesh, Daniel S. 2013. "Six Decades of Top Economics Publishing: Who and How?" *Journal of Economic Literature* 51 (1): 162–72.
Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1): 153–61.
Hicks, John R. 1939. *Value and Capital*. Oxford: Clarendon Press.
Hirschman, Albert O. 1958. *The Strategy of Economic Development*. New Haven, Conn.: Yale University Press.
Hotelling, Harold. 1931. "The Economics of Exhaustible Resources." *Journal of Political Economy* 39 (2): 137–75.
Huet, Pauline. 2015. "Émergence et structuration de l'économie des changements climatiques (1975–2013): Analyse socio-historique d'un nouveau domaine de recherche." PhD diss., Université du Québec à Montréal; Ecole des Hautes Etudes en Sciences Sociales. https://tel.archives-ouvertes.fr/tel-01112427/document.
Jensen, Michael C., and William H. Meckling. 1976. "Theory of the Firm: Managerial Behavior, Agency Costs, and Ownership Structure." *Journal of Financial Economics* 3 (4): 305–60.
Johnson, Harry G. 1958. *International Trade and Economic Growth: Studies in Pure Theory*. Cambridge, Mass.: Harvard University Press.
Johnston, John. 1972. *Econometric Methods*. 2nd ed. New York: McGraw-Hill.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–91.

Kalaitzidakis, Pantelis, Theofanis P. Mamuneas, and Thanasis Stengos. 2003. "Rankings of Academic Journals and Institutions in Economics." *Journal of the European Economic Association* 1 (6): 1346–66.

Kaplan, Norman. 1965. "The Norms of Citation Behavior: Prolegomena to the Footnote." *American Documentation* 16 (3): 179–84.

Kelly, Michael A., and Stephen Bruestle. 2011. "Trend of Subjects Published in Economics Journals, 1969–2007." *Economic Inquiry* 49 (3): 658–73.

Kessler, M. M. 1963. "Bibliographic Coupling between Scientific Papers." *American Documentation* 14 (1): 10–25.

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest, and Money*. London: Macmillan.

Kim, E. Han, Adair Morse, and Luigi Zingales. 2006. "What Has Mattered to Economics since 1970." *Journal of Economic Perspectives* 20 (4): 189–202.

Klein, Lawrence R. 1953. *A Textbook of Econometrics*. New York: Row, Peterson and Company.

Kodrzycki, Yolanda K., and Pingkang Yu. 2006. "New Approaches to Ranking Economics Journals." *Contributions in Economic Analysis & Policy* 5 (1): 1–42.

Kreuzman, Henry. 2001. "A Co-citation Analysis of Representative Authors in Philosophy: Examining the Relationship between Epistemologists and Philosophers of Science." *Scientometrics* 50 (3): 525–39.

Larsen, Peder Olesen, and Markus von Ins. 2010. "The Rate of Growth in Scientific Publication and the Decline in Coverage Provided by Science Citation Index." *Scientometrics* 84 (3): 575–603.

Louviere, Jordan J., David A. Hensher, and Joffre D. Swait. 2000. *Stated Choice Methods: Analysis and Applications*. Cambridge: Cambridge University Press.

Lucas, Robert E., Jr. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63 (3): 326–34.

———. 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics* 22 (1): 3–42.

Maddala, G. S. 1983. *Limited-Dependent and Qualitative Variables in Econometrics*. Cambridge: Cambridge University Press.

Markowitz, Harry M. 1959. *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley.

McFadden, Daniel. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers of Econometrics*, edited by P. Zarembka, 105–42. New York: Academic Press.

McKinnon, R. I. 1973. *Money and Capital in Economic Development*. Washington, D.C.: The Brookings Institution.

Meadows, Arthur Jack. 1974. *Communication in Science*. London: Butterworths.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–1725.

Merton, Robert K. 1968. *Social Theory and Social Structure*. New York: Free Press.
———. 1973. “Science and the Social Order.” In *The Sociology of Science: Theoretical and Empirical Investigations*, 250–66. Chicago: University of Chicago Press.
Morgan, Mary S., and Duo Qin. 2015. “History of Econometrics.” In *International Encyclopedia of the Social & Behavioral Sciences*, edited by James D. Wright, 8–11. 2nd ed. Oxford: Elsevier.
Musgrave, Richard A. 1959. *The Theory of Public Finance*. New York: McGraw-Hill.
Muth, John F. 1961. “Rational Expectations and the Theory of Price Movements.” *Econometrica* 29 (3): 315–35.
Newman, M. E. J. 2010. *Networks: An Introduction*. Oxford: Oxford University Press.
Newman, M. E. J., and M. Girvan. 2004. “Finding and Evaluating Community Structure in Networks.” *Physical Review E* 69 (2): 026113.
North, Douglass C. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge: Cambridge University Press.
Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, Mass.: Harvard University Press.
Patinkin, Don. 1965. *Money, Interest, and Prices*. New York: Harper & Row.
Rayson, Paul, and Roger Garside. 2000. “Comparing Corpora Using Frequency Profiling.” In *Proceedings of the Workshop on Comparing Corpora*, 1–6. ‘Stroudsburg, Penn.: Association for Computational Linguistics.
Ritzberger, Klaus. 2008. “A Ranking of Journals in Economics and Related Fields.” *German Economic Review* 9 (4): 402–30.
Romer, Paul M. 1986. “Increasing Returns and Long-Run Growth.” *Journal of Political Economy* 94 (5): 1002–37.
Rosen, Sherwin. 1974. “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition.” *Journal of Political Economy* 82 (1): 34–55.
Rostow, W. W. 1960. *The Stages of Economic Growth: A Non-Communist Manifesto*. Cambridge: Cambridge University Press.
Ruggles, Nancy D., ed. 1970. *The Behavioral and Social Sciences Survey: Economics*. Washington, D.C.: National Academy of Sciences.
Salton, Gerard, and Michael J. McGill. 1983. *Introduction to Modern Information Retrieval*. New York: McGraw-Hill.
Samuelson, Paul A. 1947. *Foundations of Economic Analysis*. Cambridge, Mass.: Harvard University Press.
Scherer, Frederic M. 1970. *Industrial Market Structure and Economic Performance*. Chicago: Rand McNally & Company.
Schumpeter, Joseph A. 1954. *History of Economic Analysis*. New York: Oxford University Press.
Silva, Ester G., and Aurora A. C. Teixeira. 2008. “Surveying Structural Change: Seminal Contributions and a Bibliometric Account.” *Structural Change and Economic Dynamics* 19 (4): 273–300.

Silva, Sandra Tavares, and Aurora A. C. Teixeira. 2009. “On the Divergence of Evolutionary Research Paths in the Past 50 Years: A Comprehensive Bibliometric Account.” *Journal of Evolutionary Economics* 19 (5): 605–42.

Small, Henry, and Belver C. Griffith. 1974. “The Structure of Scientific Literatures I: Identifying and Graphing Specialties.” *Science Studies* 4 (1): 17–40.

Smith, Adam. (1776) 1904. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Edited by Edwin Cannan. London: Methuen.

Solow, Robert M. 1957. “Technical Change and the Aggregate Production Function.” *Review of Economics and Statistics* 39 (3): 312–20.

Sraffa, Piero. 1960. *Production of Commodities by Means of Commodities: Prelude to a Critique of Economic Theory*. Cambridge: Cambridge University Press.

Stern, David I. 2013. “Uncertainty Measures for Economics Journal Impact Factors.” *Journal of Economic Literature* 51 (1): 173–89.

Sullivan, Daniel, D. Hywel White, and Edward J. Barboni. 1977. “Co-citation Analyses of Science: An Evaluation.” *Social Studies of Science* 7 (2): 223–40.

Taylor, John B. 1993. “Discretion versus Policy Rules in Practice.” *Carnegie-Rochester Conference Series on Public Policy* 39:195–214.

Theil, Henri. (1958) 1971. *Principles of Econometrics*. New York: Wiley.

Thomas, Brinley. 1954. *Migration and Economic Growth: A Study of Great Britain and the Atlantic Economy*. Cambridge: Cambridge University Press.

Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge, Mass.: MIT Press.

Train, Kenneth. 2003. *Discrete Choice Methods with Simulation*. Cambridge: Cambridge University Press.

Wallace, Matthew L., Yves Gingras, and Russell Duhon. 2009. “A New Approach for Detecting Scientific Specialties from Raw Cocitation Networks.” *JASIST* 60 (2): 240–46.

White, Halbert. 1980. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica* 48 (4): 817–38.

Whitley, Richard. 1976. “Umbrella and Polytheistic Scientific Disciplines and Their Elites.” *Social Studies of Science* 6 (3–4): 471–97.

———. 2000. *The Intellectual and Social Organization of the Sciences*. Oxford: Oxford University Press.

Williamson, Oliver E. 1975. *Markets and Hierarchies*. New York: Free Press.

———. 1985. *The Economic Institutions of Capitalism*. New York: Free Press.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, N.J.: Princeton University Press.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, Mass.: MIT Press.

Wray, K. Brad. 2005. “Rethinking Scientific Specialization.” *Social Studies of Science* 35 (1): 151–64.

Zimmermann, Christian. 2013. “Academic Rankings with RePEc.” *Econometrics* 1 (3): 249–80.