



The two sides of the Environmental Kuznets Curve: a socio-semantic analysis

Telmo Menezes, Antonin Pottier, Camille Roth

► To cite this version:

Telmo Menezes, Antonin Pottier, Camille Roth. The two sides of the Environmental Kuznets Curve: a socio-semantic analysis. *Œconomia - History/Methodology/Philosophy*, 2023, 75 (13-2), pp.279-321. 10.4000/oeconomia.15729 . hal-04319048

HAL Id: hal-04319048

<https://hal.science/hal-04319048>

Submitted on 2 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The two sides of the Environmental Kuznets Curve: a socio-semantic analysis

Telmo Menezes^{*1}, Antonin Pottier^{†1,2,3}, and Camille Roth^{‡1,4}

¹Centre Marc Bloch Berlin (CNRS/HU), Friedrichstr. 191, 10117 Berlin, Germany

²CIREN (Centre international de recherche sur l'environnement et le développement), EHESS, Paris, France

³Wissenschaftskolleg, Berlin, Germany

⁴CAMS (Centre Analyse et Mathématique Sociales, UMR 8557 CNRS/EHESS), 54 Bd Raspail, 75006 Paris, France

Abstract

Since the 1990s, the Environmental Kuznets Curve (EKC) hypothesis posits an inverted U-shaped relationship between pollutants and economic development. The hypothesis has attracted a lot of research. We provide here a review of more than 2000 articles that have been published on the EKC. We aim at mapping the development of this specialized research, both in term of actors and of content, and to trace the transformation it has undergone from its beginning to the present. To that end, we combine traditional bibliometric analysis and semantic analysis with a novel method, that enables us to recover the type of pollutants that are studied and the empirical claims made on EKC (whether the hypothesis is invalidated or not). We principally exhibit the existence of a few epistemic communities that are related to distinct time periods, topics and, to some extent, proportion of positive results on EKC.

Keywords: Environmental Kuznets Curve (EKC), Socio-semantic networks, Semantic hypergraphs, Bibliometric analysis, Citation networks

The relationship between environmental impacts and economic activity is a hotly debated topic in environmental economics since its inception in the 1960s. At that time, harsh criticisms against economic growth were released as economic expansion was considered the source of environmental degradation. The summit of this wave was the *Limits to growth* – Meadows et al. (1972) report, which is famous for having

^{*}menezes@cmb.hu-berlin.de

[†]pottier@centre-cired.fr

[‡]camille.roth@cnrs.fr

made the case that economic growth was not sustainable, either because of the lack of raw resources it constantly demands or because of the amount of pollution it unleashes. This call to put growth to an halt raised strong rebuttals among economists who contented that economic progress and substitution can accommodate the economy with any resource scarcity. Fears of expanding pollution were however less easily dismissed. Hence an already famous environmental economist like William Nordhaus could write in 1992 that “long-run constraints upon economic growth might well exist [...] one possibility is that the scale of human activity could overwhelm the capacity of the globe to tolerate industrial wastes” (Nordhaus, 1992, 3-4).

Concerns about the need to combine development and environmental protection had led in the meantime to formulating the concept of *sustainable development* and making it a goal for the international community. Sustainable development means “meeting the needs and aspirations of the present without compromising the ability to meet those of the future”, as was stated by the Brundtland et al. (1987, chap. 1, §49) report, which added that, “far from requiring the cessation of economic growth, it recognizes that the problems of poverty and underdevelopment cannot be solved unless we have a new era of growth”. How this new era of growth could be achieved without degrading further the environment was however left unaddressed.

In 1992, the *World Development Report* suggested a reassuring way out of the dilemma (World Bank, 1992). It displayed several patterns of environmental indicators depending on country income: monotonically decreasing, monotonically increasing, or increasing then decreasing. According to the latter pattern, the state of the environment, after an initial period of degradation, will automatically improve. The promise was clear: the economic development would somehow take care of its own drawbacks. Environmental degradation would not so much be an issue contrary to accelerating economic growth, as the latter would cure the former. As Beckerman (1992, 491) put it, “in the longer run, the surest way to improve your environment is to become rich”.

This inverted-U relationship between economic output and environmental degradation was soon named the environmental Kuznets curve (EKC) hypothesis (Panayotou, 1993), from the Kuznets curve, an inverted U-shaped relationship between inequality and economic development that Simon Kuznets (1955), one of the first national accountants, had observed¹. According to the EKC hypothesis, environmental degradation has an inverted U-shaped relationship with economic development. Many empirical contributions immediately investigated whether such a relationship exists or not, for what type of pollutants, while others pinpointed the problems in estimating its parameters. A literature subfield, both theoretical and empirical, grew quickly around the “environmental Kuznets curve”.

Thirty years after these seminal contributions (Grossman and Krueger, 1991; Shafik and Bandyopadhyay, 1992; Panayotou, 1993) more than 2000 articles belong to this subfield. To assess what this literature has finally to say on the validation or not on the EKC hypothesis is challenging, given its size. This is a problem faced by many social scientists: the significant growth in published articles makes it harder to identify, analyze, and assess trends, topics or influential ideas in a given field.

The difficulty of the task motivates the research and the development of an auto-

¹Ironically, this Kuznets curve is no more than a fortuitous relation, see Piketty (2014, introduction).

mated analysis, that we present here in this special issue on “quantitative and computational methods in the social studies of economics”. We hope to foster the application of computational techniques to a broader range of subjects and to stimulate the discussion around their possibilities.

The purpose of this article is first to describe the evolution of the EKC literature. To that end, we quantitatively investigate this literature with tools from social analysis and semantic analysis, hence leading to a socio-semantic analysis. Social analysis refers to the study of networks between actors: here we used the various bibliometric networks that can be derived from the oriented relation that a citation creates between two articles. Semantic analysis refers to the automatic study of the meaning of sentences – in our case by detecting linguistic patterns that we identified for this specific application – and it will be applied to the abstracts of the collected articles.

We aim at assessing how the empirical literature of the EKC has developed and whether it has come to a conclusion. However, such an assessment is more complicated than it may seem. Indeed, the literature on EKC has gone in many different directions and has undertaken significant changes since the first contributions, to the point that validating the EKC hypothesis has no longer a univocal meaning. We therefore also map the development of this specialized research, both in term of actors and of content, and trace the transformation it has undergone from its beginning to the present.

The success of the EKC concept that originates within economics, its spreading in several related disciplines and its evolution offer however broader lessons that go beyond this particular subfield. Our study shows how a concept like the EKC has actually become a keyword on which research from various backgrounds hangs on. This evolution could be seen as another example of economic imperialism but it is actually more accurately described as a combination of attraction of the EKC concept on related disciplines and outpouring of the EKC research in other disciplines. In the process, the EKC research field has lost part of its initial cohesiveness: it has been, in a way, a victim of its own success. Our study thus provides with the EKC an interesting case in the social studies of economics to understand the diffusion of economic concepts outside economic literature.

Finally, our article also innovates on the methodological side by combining semantic and socio-bibliometric analysis. Social network analysis has been advocated to be used in the history of economics, for example to study collaboration and interactions in scientific communities (Claveau and Herfeld, 2018; Herfeld and Malte, 2018). We use a social network derived from bibliometric data and we supplement bibliometric analysis with an automatic recognition of the claims made by the article. This allows us to recover what an article does or claims to do and thus to map the social relations among actors with the actual content of the scientific disputes between them. The approach developed in our study can be readily used to investigate other fields or concepts. The case study on the EKC research allows us to reflect on the advantages and limitations of the approach.

The article is organized as follows. Section 1 reviews the literature on the subfield of the EKC and positions this article among the bibliometric analyses of this subfield. Section 2 presents the methodology of semantic analysis and the first results in terms of temporal development of the field, type of pollutants that are investigated, and type of claims that are made. Section 3 introduces the analysis of the citation network within

the corpus, contrasts its various blocks and the dominant individuals and journals, while discussing its temporal features and relying on the instruments provided by the previous section. Section 4 furthermore discusses limitations of our analysis, followed by a summary of our findings in Section 5.

1 Early history of the Environmental Kuznets curve and literature review

The environmental Kuznets curve hypothesis encapsulates the idea that environmental impacts increase then decrease along with economic development. The modern formulation of this idea is the product of several studies conducted simultaneously in the early 1990s. It can be traced back to Grossman and Krueger (1991), who studied the environmental impacts of the free-trade agreement between the USA, Canada and Mexico. They observed that the ambient concentrations of some pollutants (SO_2 or particulate matter) increase at low income but decrease at high income, hence suggesting that trade liberalization that would foster economic development could also be good for the environment. The same idea also inspired Shafik and Bandyopadhyay (1992) who studied more systematically the link between economic growth and environmental quality, considering eight indicators of environmental quality across a large number of countries. This research was given wide outreach as it was a background article for the *World Development Report 1992* (World Bank, 1992), which incorporated some of its finding as we have explained in the introduction.

After the publication of the *World Development Report*, the number of publications concerned with the inverted U-shaped relationship between the environmental degradation (measured by several indicators of pollutants) and economic development (measured by GDP or GDP per capita) increased quickly and used the phrase “Environmental Kuznets curve” coined by Panayotou (1993) at the International Labour Office to refer to this particular relationship. From the start, the EKC hypothesis triggered heated debates. Empirical tests, based on cross-section of countries or panel data, gave varying results, depending on the type of pollutants (water and local air pollution were mostly investigated), the sample of countries or the time period. The actual significance of the EKC was also contested: most articles were interested in the value of the turning point, that is the income after which pollution would decrease, yet Cropper and Griffiths (1994) argued that the height reached by the EKC at its peak was just as important. The political implications of the existence or not of an EKC were no less controversial (Panayotou, 1997). Along the lines suggested by Grossman and Krueger (1991), the mechanisms that would explain such a relationship were also discussed, distinguishing the effects of an expanding scale of the economy, of composition change of output, and of the progress of technologies.

In 1997, a special issue in *Environment and Development Economics* attempted to provide new insights on this growing field. Barbier (1997) in its introduction summed up the main directions taken by the research on EKC. He synthesized the results obtained for a large varieties of pollutants (classified by broad categories: air pollution, water pollution, deforestation and others) and discuss the level of income at which the

turning point occurs. He already pointed to features of the field that will be long lasting, i.e. that the EKC hypothesis remains highly contested and controversial.

The potential for a proliferation of articles on EKC can already be discerned there. Barbier (1997) indeed remarked that the EKC is a falsifiable hypothesis and thus lends itself to be tested in various national contexts, with different methodologies, for several pollutants, etc. Because it was formulated in a falsifiable form, yet sufficiently vague (which pollutants? time frame? countries?), one could see in retrospect that it calls for numerous, and possibly endless, empirical studies.

Explaining EKC also raised a theoretical challenge. Some theoretical models were indeed proposed (Andreoni and Levinson, 2001; Brock and Taylor, 2010; Kijima et al., 2010), but they have not been produced on the same scale as the empirical studies (see Kijima et al. (2010) for a survey).

Surveys were necessary to keep track of the development of the research field and they became more and more dense to address its many branches. Dinda (2004) surveyed the literature, which he already found “sizeable”. He discussed the empirical estimations and theoretical explanations, as well as the conceptual and methodological critiques raised against the EKC. His survey was 25 pages long and included around 180 references.

The more recent survey by Shahbaz and Sinha (2019) focused on empirical articles dealing specifically with CO₂ emissions i.e., almost 200 articles published before 2017. Although they restrict themselves to a single type of pollutant, they find the research to be inconclusive. The results are wildly varying, which they relate, as in the early days of the field, to the many methodologies, contexts, time frames, data sets that the researchers can choose.

Aside surveys directly engaging with the EKC hypothesis and synthesizing results for researchers of the field, several articles examined this field from an external perspective, using bibliometric tools. Sarkodie and Strezov (2019) has done a meta-analysis aiming at identifying the value of the turning point found by studies validating the EKC hypothesis (though irrespective of the type of pollutants). They also offer a section with a bibliometric analysis of the EKC research. They mainly investigate author contributions, citations, journals, countries, as well as topics. Their analysis is however static and cannot be used to map the evolution of the field. Bashir et al. (2021) investigate the citation network of the EKC literature. They compute basic statistics, like the countries or research institutes contributing most to the literature, the journals in which the literature is published, the most cited studies. They also draw collaboration networks between authors and look for trending topics within the EKC literature. Koondhar et al. (2021) investigate the co-citation network within the EKC field with the “CiteSpace” software. They focus on relations between authors, journals, institutions and keywords. Anwar et al. (2022) have similarly done a bibliometric review from data of the so-called Web of Science database. They also use the “CiteSpace” software to identify key articles according to various provided metrics. However their results are not interpreted: they simply list the articles found and sum them up briefly so that as it stands their work does not allow an understanding of the development of the EKC literature. Finally, Naveed et al. (2022) focus on the most influential journals, articles and authors, relying on citations analysis of articles indexed in the “Scopus” database. They also analyze and classify the content of the 100 most influential articles.

These studies of varying depth however do not address what we want to investigate. First, the analysis relies on keywords to discuss the content of the article, thus ignoring its meaning, that we can catch with semantic analysis. Second, they are mainly static and they lack an historical perspective, whereas we map the development of the field through time. Third, they do not pay attention to the transfers that occurred across disciplines and journals, which we highlight in our analysis.

2 Methods and descriptive results

2.1 Corpus building

We build our corpus from the Scopus database. Scopus is one of the main bibliographic databases, released by the scientific publisher Elsevier.² We query Scopus for articles containing the exact phrase “environmental Kuznets curve” in the title, abstract or keywords. The search was performed on June 4, 2021 delivering 2709 articles. We collect all possible data that were furnished by Scopus. Here, the analysis will mainly be focused on titles, authors, abstracts and citations. We screen random references to see whether the collected articles were mis-attributed. Given that the syntagm “Environmental Kuznets curve” is already quite narrow, we find no article that addresses a different topic, so we do not dwell further to exclude wrongly included articles, that is we retain the whole database as the corpus of our analysis.

Figure 1 shows the number of articles in our corpus per year of publication. It shows a strong increasing trend, as 54% of the articles have been published in the last four years and a half (since 2017 included), although the first articles are from 1994 (only 2% have been published before 2000 included). We find that in the last decade the number of article grows at a 20% rate per annum, far more than the growth rate of between 5% and 6% found by Claveau and Gingras (2016) for articles in economics.

We consider three periods of time: the last two five-year periods (2012-16 and 2017-21), and the beginning of the dataset (1995-2011) gathered in a single period for the sake of significance (there are comparatively very few articles in these early stages, as shown in figure 1), to strike a sensible balance between article density and temporal resolution. As a result, the 1995-2011 period comprises 617 articles; 2012-2016, 621 articles, and 2017-2021, 1467 articles. There are 4435 unique authors in the corpus.

Scopus gives for each article of the corpus the list of references of that article. Parsing this list of references, identifying each and finally linking it back to an article of our corpus enables to build the citation network within our corpus. All citation analysis will be based on this network and its derivatives. Hence, all metrics referring to number of citations come from this and so have to be understood as internal i.e., most-cited articles are articles that are most cited by the articles of the corpus. These pre-processing steps were performed using simple Python scripts that we developed for this purpose.

²<https://www.elsevier.com/solutions/scopus/how-scopus-works>




Figure 1: Corpus per year of publication

2.2 Semantic analysis and semantic hypergraphs

We develop an automatic pattern recognition approach to extract claims about the EKC from abstracts — in a nutshell, we are able to classify the recognized claims as positive and negative results on the EKC, and the variables they are based upon. It is especially interesting for an analysis of the field as it allows to reach, to some extent, what an article claims, and in turn to deepen traditional bibliometric analysis by complementing it with semantic analysis. The same abstract may report both positive and negative results. Since our method works at the level of individual claims, we can and do detect such mixed cases. A limitation inherent to the corpus we use is that only abstracts are considered, and sometimes negative results are reported only in the main text of the article.

Our approach is based on *Semantic Hypergraphs (SH)*, a knowledge representation model that is intrinsically recursive and accommodates the natural hierarchical richness of natural language. A full description of SH can be found in Menezes and Roth (2019). SH have been recently proposed, but they have already been used for example in the study of the credibility of research impact statements (Bonaccorsi et al., 2021).

One popular contemporary alternative would be to train a classifier with the help of a large language model, e.g. of the BERT family. A disadvantage of this approach is that, given its purely statistic nature, it typically require large training datasets. This disadvantage is accentuated by the fact that this NLP task is somewhat delicate, particularly in identifying if a result is positive or negative, and understanding negation appears to be a typical shortcoming of current deep learning models (Ettinger, 2020). As will become clear in the subsequent paragraphs, identifying if a result refers to a confirmation or refutation of EKC also requires some nuance. SH makes it possible to infer a useful classifier from a small number of cases – which is not only important because of the human effort required to annotate training datasets, but perhaps even

crucial because our dataset is relatively small.

We give an overview of SH in Appendix A. In a nutshell, SH makes it possible to define and extract relevant linguistic patterns in a semi-supervised manner. More precisely, focusing on a subset of 500 randomly selected sentences, we first extract from the most frequent predicates the ones that are associated with result claims. We add conditions to focus on results related to EKC, including rules to capture the notions of "U-curves" and "N-curves" which, in the corpus, implicitly refer to EKC. We finally take negations into account in order to distinguish between positive and negative results.

On the whole, we design a handful of patterns which are able to classify sentences featuring an EKC claim, and whether it denotes a positive or a negative result. To evaluate the performance of our semantic classifier, we randomly selected 3 sets of 50 abstracts each, to be manually annotated by the 3 authors, so that each author annotated one of the sets. Notice that abstracts can present no results, only positive results, only negative results, or both positive and negative results. Annotation was performed with no knowledge of the automatic classification. Overall precision and recall ³ results are shown in table 1.

<i>EKC validation claim type</i>	Precision	Recall
<i>positive result</i>	.809	.847
<i>negative result</i>	.833	.366

Table 1: Precision and recall of EKC result classifier when compared against a randomly-selected set of 150 manually-annotated articles.

We found both precision and recall to be satisfactory for positive claims. For negative claims, precision is satisfactory but recall is poor, which is to say that our classifier markedly underestimates the number of negative claims. We found this to be related to a tendency by authors to present negative results in a less explicit fashion: negative results are often expressed with more convoluted sentences, and with a lot of qualifications. Often, negative results are also implied and thus cannot be identified through direct claims. A telling example is Acaravci and Ozturk (2010), a largely cited article (155 times, ranking 7th in the corpus), which is automatically classified as claiming a positive result because the abstract indicates that “These results support that [sic] the validity of environmental Kuznets curve (EKC) hypothesis in Denmark and Italy”. However, the article investigates the relationship between CO₂ and real GDP per capita for 19 European countries: the abstract thus implicitly implies that a relevant EKC relationship was *not found* for the remaining 17 countries, which we confirmed by examining the article.

The much larger recall error of the algorithm for negative results leads us to correct

³*Precision and recall* are common measures in Machine Learning, used to evaluate a classifier’s performance. Precision is the fraction of true positives out of all positive predictions, while recall is the fraction of true positives out of all actual positive observations in the data.




Figure 2: Estimated percentage of articles with a positive or with a negative result.

all values found by the algorithm by multiplying raw quantities of results by the associated precision and dividing them by the associated recall.⁴ For instance, if a negative result is detected in N articles, we consider that there are actually $\frac{N \cdot 0.833}{0.366}$ such articles. Notice that negative claim precision also contributes to positive claim precision: even if negative recall is relatively weak, negative precision is still quite useful because it prevents claims from being wrongly classified as positive claims.

In figure 2 we plot the ratio of articles with at least one positive result and the ratio of those with at least one negative result, knowing that a small portion of articles feature both. We see that the proportion of articles with positive results and that of articles with negative results do not exhibit a very marked difference. Even though there has been a modest albeit increasing lead in the number of articles with positive findings over the recent years, there nonetheless remains a substantial proportion of articles with negative results. From this analysis of the meaning of the abstracts collected in our corpus, we can confirm that the literature on EKC is generally inconclusive, as stated by several observers or actors of the field (Shahbaz and Sinha, 2019; Haberl et al., 2020). In other words, the evidence cannot be said to be settled and there is an on-going controversy, which may reflect different scientific practices and epistemic communities. We will comment more on that in section 4.

2.3 Topics

The EKC is a generic concept that can be applied to the relationship between an economic development variable (such as GDP) and various pollutants and environmental

⁴Precision corresponds to the proportion of correctly detected items of a certain category among those detected as belonging to that category, while recall is the proportion correctly detected items among those belonging to that category. Hence, precision/recall is the rate of correctly detected items per items detected as belonging to a category (for more context on such correction factor see e.g. Schmid et al., 2020).




Figure 3: Percentage of articles evoking a relationship on a given topic (sums over 100% as an article may address several topics, bar elements representing less than 4% were not labeled).

pressures. We now identify which kind of environmental variables have been investigated in a given article.

Once again we take advantage of SH representation. We found that the pattern “relationship between X and Y” is prevalent across the corpus, and that it largely corresponds to a relationship between an economic variable and an environmental variable.⁵ We exploited it to identify all the environmental variables present in X or Y. The most common ones were manually grouped into eight categories: greenhouse gases (labeled as “GHG”, which includes for a very large part CO₂-related research), energy, water, waste, environmental footprint, sulphur oxides (SOx), nitrogen oxides (NOx) and local air pollutants (like particulate matter or carbon monoxide, this category does not include SOx nor NOx). We treat these environmental variable categories as topics.

On the whole, 82% of articles were found to contain at least one of these environmental variable categories. On figure 3 we present the ratios of articles mentioning each category. Ratios sum to more than 100% since, although topics are non overlapping, an article can be classified in several categories: an article can analyze the EKC both for air and water pollutants. Energy and its various vectors (oil, coal, gas) were frequently associated with income, yet it appears that articles are often both categorized as energy and carbon. This is because the articles that investigated the link between GHG emissions and income often add various form of energy as control variables.

This analysis yields three main results. First, the strongly increasing focus on GHG/CO₂ emissions and energy, which was also found by Sarkodie and Strezov (2019)

⁵The pattern can be represented in SH as: (between/B relationship/C (and/J X Y)).

using a different method (keyword analysis). This focus contrasts with what one might have anticipated based on the early developments in the field. Indeed, although Shafik and Bandyopadhyay (1992) considered CO₂ within their set of pollutants, the early decade of research on EKC was not strongly concerned with CO₂. Dinda (2004)'s review mentions research on CO₂ but focuses on local air pollution and water pollution. One could therefore have expected that the literature would go on focusing on these pollutants, as the EKC hypothesis was heavily debated for these ones (see for example Stern and Common (2001) on sulfur emissions). Furthermore, as Dinda (2004, p. 441, 449) makes clear, the literature at that time did not find the evidence of EKC for global pollutants like CO₂. On the contrary, from 2014, the articles identified as investigating CO₂ or GHG constitute more than half of our corpus, representing above 55% thereof in recent years. The trend is quite massive and shows no sign of losing steam. Second, while seminal publications discussed various pollutants, especially air pollutants such as nitrogen and sulfur oxides (NO_x and SO_x), and various environmental stressors such as waste, more recent publications do not seem to focus explicitly on such specific pollutants: the shares of SO_x, NO_x, waste as well as water [pollution] strongly decreased; only the share of the generic reference to "local air pollutants" remained stable. To summarize, we witness a strong turn to energy and GHG-related pollutants from around 2012.

This turn can be explained. First and foremost, the growing number of articles targeting GHG emissions has to be related to the increased salience of climate change as the environmental problem vital for human societies. This means that researchers are more likely to focus on GHG and neglect other environmental problems, but also that research is more easy on this topic because of the large availability of dedicated databases (on GHG emissions) and of the ease of attracting fundings. Furthermore, methodological changes are also involved in this turn. Ang (2007) was the first to bring together a literature focused on the relationship between economic output and energy, and the EKC literature, hitherto focused on the relationship between economic output and pollutants (here specifically CO₂ emissions). This promoted the integration into the EKC literature of studies considering both energy and GHG emissions, while transforming the EKC itself, as it became a relationship between economic output and GHG emissions within a broader framework. We come back to this transformation in the discussion (see section 4).

A cross analysis of positive/negative results vs. topics makes it possible to refine the picture further. Table 2 provides the relative proportion of positive vs. negative results for each topic through time — in other words, it paints the categories of figure 3 with the categories of figure 2. We see that the literature of the last period 2017-21 features a majority of positive results (proportion > 1) for all topics. This reflects the aggregate trend previously observed for that period, for which the share of positive results significantly increases (to reach 53% vs 38%). This trend is however not uniform: results related to footprint, GHG, local air pollutants, NO_x and, to a smaller extent, SO_x, experience the strongest increase in positive results. More interestingly, for half of the categories (local air pollutants, footprint, water and NO_x), trends are reverted from a majority of negative results to a minority; the same could almost be said for GHG (from half-half to strongly positive-leaning results). Waste is the only topic that exhibits a decrease in the proportion of positive results from the first to the final period.

<i>Period</i>	Topic							
	GHG	Energy	Local air pollutants	water	SOx	Waste	Footprint	NOx
<i>1995-2011</i>	1.03	1.29	0.84	0.81	1.18	1.10	0.31	0.91
<i>2012-2016</i>	1.05	1.09	1.15	1.03	1.00	0.67	0.42	0.92
<i>2017-2021</i>	1.55	1.39	1.64	1.01	1.39	1.01	1.07	1.50
<i>aggregate</i>	<i>1.35</i>	<i>1.31</i>	<i>1.32</i>	<i>0.95</i>	<i>1.19</i>	<i>0.94</i>	<i>0.88</i>	<i>1.14</i>

Table 2: Evolution of the ratio of articles with positive vs. negative results, broken down by topic, computed over the three defined time periods (figures in bold indicate for each topic the period of maximum ratio) and for all periods (“aggregate”).

Thus, while the increase in topic-related research suggests that the field is getting more specialized or, at least, giving more attention to specific relationships, this tendency also comes with an evolution of the imbalance around the confirmation or refutation of EKC that seems to affect more certain topics than others. We shall see below how this topical specialization is distributed on the author network.

3 Citation network and blocks

To have a better sense of how the field has developed and why, we now turn to the analysis of the author citation network.

To begin, we compute for each author the number of times their articles have been cited. As said before, we only consider citations internal to the corpus. This means that the citations attracted reflect the influence of the articles in the field of the EKC research. With this, we cannot assess the impact of the articles of the corpus in the environmental economics literature in general, nor are we able to assess to the influence of articles outside the corpus on the corpus.

We concentrate on the most (internally) cited authors, listed in table 3, for whom we can already distinguish several patterns. I. Öztürk and M. Shahbaz entered the fields in the 2010s, and since then they have published a large number of articles. D.I. Stern has been present in the field over the whole time span, with a large yet lower number of articles per year. Soumyananda Dinda has been mainly influential through his review of 2004 published in *Ecological Economics* (673 out of 879 citations).

The table is a little bit different from the one obtained by Sarkodie and Strezov (2019). First two authors are the same, although in reverse order. Shahbaz and Dinda are also present in the top-five most cited authors of Sarkodie and Strezov (2019). The differences seem to stem from three facts. First Sarkodie and Strezov (2019) use the *Web of Science* database to extract their corpus whereas we rely on *Scopus*. Second, they use metadata citation counts that apply to the whole Web of Science database and thus are likely to include citations outside of the EKC research field. Third, our analysis is done three years later. As Öztürk is a prolific and active author (as we will see), this may explain why he has overcome Stern in the mean time.

Author	#A	c	Activity time span
Öztürk I.	35	1779	2010-2021
Stern D.I.	17	1622	1996-2020
Shahbaz M.	49	1524	2012-2021
Dinda S.	7	879	2000-2009
Al-Mulali U.	13	756	2015-2019

Table 3: Five most cited authors. #A is the number of articles, c number of citations.

3.1 Blocks of authors

We now look into more structural characteristics of the citation network. To do so, we extracted the citation network between authors from the corpus from the citations between articles. Citations are matched to articles in the corpus if both their titles (case-insensitive) and years of publication are the same. Each citing article induces edges of weight 1 from all its authors to each author of the cited article. This is a directed weighted network, with each edge weight indicating the number of times the origin cites the target. This network is very dense, so to make it tractable we adopt a simple edge removal and graph simplification strategy.⁶ We consider only 3 outgoing edges per author, pointing at the 3 most cited target authors by each actor of origin. This means that every node has out-degree 3 and any number of incoming edges (in-degree). Furthermore, we consider only authors that received at least 10 citations. We have found these thresholds to be remarkably robust. For example, employing a threshold an order of magnitude higher (top 30 most cited targets) produces very similar block partitions. To also be able to observe the temporal evolution of the network, we compute cumulative networks for three time periods starting in 1995 and ending, respectively, in 2011, 2016 and 2021, the last one corresponding to the full dataset.

We then use Stochastic Block Models (SBM) (Holland et al., 1983; Peixoto, 2018) and the open-source library *graph-tool* to infer the structure of the network⁷. SBM is a generative model, where each node is assigned to a given block. SBM blocks are such that nodes of a given block connect to nodes of another block (that can be the same)

⁶A myriad of graph simplification methods (variously denoted as graph sampling, reduction, compression) have been proposed by the network science community over the last two decades, principally aimed at easing structure visualization or statistical feature computation (Zhao et al., 2020). Node-centric edge sampling (such as favoring high-degree nodes e.g., Ribeiro and Towsley, 2010) and network-level edge selection (such as minimal spanning trees e.g., Van Ham and Wattenberg, 2008) are typical strategies fulfilling relatively distinct goals: while the former seem to operate on local structure, the latter rather help in preserving a sense of global connectivity (Jalali et al., 2016), even though they can do both (respectively Leskovec and Faloutsos, 2006; Jia et al., 2008). As we are interested in blocks we rather focus on the former and adopt a very simple approach of uniformly pruning outgoing edges in order to concentrate on the most important targets, while preserving the structure of incoming edges i.e., received authority from peers – admittedly a key feature of the citation networks we study. While much simpler, it is close in spirit to node-centric sampling methods which tends to favor links to higher degree nodes and demonstrate satisfying preservation of community structure (Maiya and Berger-Wolf, 2011; Ahmed et al., 2011; Voudigari et al., 2016).

⁷<https://graph-tool.skewed.de>

with a constant probability. SBM jointly infers block membership and the matrix of connection probabilities, such that the likelihood of the observed network given the generative model is maximized. Put differently, blocks represent similar connection patterns from the nodes of that block towards nodes from other blocks. One feature of SBM is that maximizing this likelihood also implies a minimization of the amount of information necessary to describe the model (description length). In other words, there is a preference akin to *Occam’s Razor* for the simpler model. We performed model selection based on description length and opted for *degree-corrected* SBM (Karrer and Newman, 2011), since it achieved significantly lower description lengths in comparison to standard SBM. Given that SBM inference is a stochastic non-deterministic process, we performed it 10K times and selected the one that achieved lowest description length. We arrived at a partition in 4 blocks for the full network covering the entire time period under study, as depicted in figure 4(c).

Blocks are distinguished by significant linking behavior differences, both inwards and outwards, which can be interpreted as an indication of their role in the network. For example, block A contains a small number of highly cited authors. A tends to be cited by all blocks and to not cite any other blocks, while all others tend to cite A.

Figure 4(c) shows us that two large blocks (B and C), of comparable size, account for the majority of authors in the network. Given that the two large blocks are not connected in figure 4(d), this could suggest the existence of two separate factions with the similar internal dynamics of a large “following” and a smaller number of highly influential authors. Notably, Stern belongs to block A and Öztürk to block C. While block A is clearly influential to all the other blocks, its topological insertion suggests its closer alignment with block B. The remaining block (D) is dominated by Chinese authors publishing articles with a strong focus on China.

Notice that figure 4(d) is rendered with a threshold of probability of connection of $p > 0.01$, and that there is some probability of connection between all pairs of blocks that is greater than zero, and that such connections do exist. Furthermore, the formation of such connections is affected by size effects, which is why even with low affinity below the 0.01 threshold, blocks B and C still have a good amount of connections between them. It is also affected by age: an older article had more time to be cited and can only cite articles that were published before it.

In figure 4(d) we can observe another interesting topological fact: although Stern has been overtaken by Öztürk in terms of number of citations received within the EKC literature, the small block of authors that he belongs to appears to remain the most central and influential. We can see that all other blocks have a strong tendency to cite A, and that A is in fact the only block *for which this is the case*. Another indication of this asymmetry is the more casual observation that Öztürk cites Stern’s articles thirteen times, whereas Stern does not cite Öztürk at all – this being true also for the full citation graph, not just the simplified one with the 3 outgoing edges per author. When interpreting this asymmetry, it is also good to keep in mind that Stern’s work is older than Öztürk’s.

In table 4 we present a set of metrics for the blocks that are not of a topological nature. This means that the distinctions that these metrics provide are not implied by the network structure, and therefore help to strengthen the hypothesis that these blocks do indeed correspond to different practices within EKC research. Blocks A and then




Figure 4: Graphs of citations from authors to authors (or blocks of authors to blocks of authors in the last panel), colored by blocks determined by SBM. Node radius is proportional to in-degree. Three periods (cumulative): (a) 1995-2011 (b) 1995-2016 and (c) 1995-2022. (d) Graph of author blocks for the last period. Edge thickness is proportional to the overall probability p of any given author from one block citing any given author in the other block. Very low probabilities of connection ($p \leq 0.01$) are not depicted. Node radius is proportional to the total number of citations (for the entire corpus) received by the authors in the block.

<i>Block</i>	<i>#Authors</i>	<i>Pos</i>	<i>Neg</i>	<i>Endogamy</i>	<i>Year</i>	<i>#A</i>	$\overline{\#A}$	<i>k</i>
<i>A</i>	3	0.354	0.084	0.003	2007.9	9.00	0.79	5.0
<i>B</i>	360	0.422	0.416	0.033	2012.1	3.11	0.76	4.9
<i>C</i>	406	0.562	0.407	0.102	2017.9	4.47	1.24	9.2
<i>D</i>	57	0.492	0.441	0.085	2016.8	5.96	0.96	15.7

Table 4: Various metrics per author block, including ratios of positive and negative EKC claims, endogamy measured as ratio of citations received from co-authors, mean year of publication, mean number of articles per author (*#A*) and further normalized per year ($\overline{\#A}$), and mean number of unique co-authors (*k*).

B have a lower ratio of positive results than C. The ratio of negative results is more similar across blocks, except for A. There is also a clear difference in endogamy, with C being the most endogamic and both A and B the least. Interestingly, block D appears to be a middle ground between A+B and C in all metrics, as well as in its topological insertion in the citation network. This highly regional block might be influenced by the two main blocks.

The temporal aspect is also of interest. Considering the mean year of publication for the articles by authors in each block, we can see that block A is the oldest, followed by B, and then C and D are much more recent. This invites the hypothesis of a shift in practices having taken place during some period in time. In fact, the three steps of the longitudinal study, corresponding to the panels (a), (b) and (c) of figure 4, reveal a complexification of the citation network from uni- to bi-polarity. Blocks in the two first networks were colored according to the block in the final network with which they have the highest Jaccard similarity. It is clear that the field was dominated by Stern and Dinda in the first period, then assigned to block B. Öztürk and Shahbaz appear in the second period along with block C, with the bipolarity already fully developed, and Stern still appearing as the most influential member of block B. In the final period, Stern and two other influential and seminal authors detach into block A and block D is formed.

Finally, the breakdown of the presence of each topic for each block, shown in table 5, also paints a quite heterogeneous picture. It indicates a certain level of focus and specialization proper to some blocks: for instance, A on SO_x, C on GHG and energy and D on local air pollutants.

3.2 Focusing on the two leaders

David I. Stern, the second most cited author, has been active in the field for 25 years. He began his career⁸ with a PhD in Geography from Boston University. From 1996, he was a Research Fellow at the Centre for Resource and Environmental Studies in Australia. He is currently Professor at the Crawford School of Public Policy, Australia.

⁸Information available at <http://sterndavidi.com>

<i>Block</i>	GHG	Energy	Local air pollutants	Water	SOx	Waste	Footprint	NOx
<i>A</i>	.333	.074	.222	.000	.333	.000	.000	.000
<i>B</i>	.345	.257	.344	.123	.109	.084	.020	.042
<i>C</i>	.664	.618	.310	.062	.039	.043	.085	.023
<i>D</i>	.436	.392	.537	.128	.084	.106	.018	.040

Table 5: Percentage of articles mentioning each topic for each author block. Bold figures indicate the block where a topic has the highest presence.

He was associate editor of *Ecological Economics* (which appears in Table 6) from 2002 to 2018 and belongs to its editorial board since then. His most cited article is “Is There an Environmental Kuznets Curve for Sulfur?” (Stern and Common, 2001) which uses panel data to investigate the EKC for sulfur emissions and essentially shows that earlier findings of an EKC can be explained by a sample restricted to high-income countries: an EKC could not be found using the global sample.

Ilhan Öztürk, the first most cited author according to our investigation, is a newer scholar who completed his PhD in Economics in 2009 from Çukurova University in Turkey and began publishing on EKC in 2010.⁹ He has been working at Çağ University since 2000, where he became a professor in 2017. Öztürk is editor-in-chief of the *International Journal of Energy Economics and Policy* (which appears in Table 6), founded in 2011 and edited by EconJournals, a platform run by Ilhan Öztürk. This publisher has two other journals: Ilhan Öztürk is also editor-in-chief of the first and co-editor of the second. On EKC, his most cited article is “Investigating the environmental Kuznets curve hypothesis in Vietnam” (Al-Mulali et al., 2015). It uses time-series of carbon emissions, GDP, and several controls (imports, exports, various forms of energy) to test for the EKC and finally rejects it.

Öztürk and Stern are the two most cited authors but, given the numbers available in table 3, they seem to be inserted differently in scientific networks. As said before, Stern has published less than Öztürk in our corpus, 17 articles against 35, but within 25 years of activity compared to 12, which corresponds respectively to 0.68 articles on EKC per year vs. 3.2. This difference is also visible in the number of unique co-authors: only 10 for Stern vs. 83 for Öztürk. These different practices extend to citations, whereby the 1808 citations of Öztürk are more concentrated than the 1644 citations of Stern, in terms of both citing authors and citing articles: 1551 distinct authors cite Öztürk vs. 1992 for Stern (hence 1.17 cites per author compared with 0.83), and 794 distinct articles for Öztürk vs. 1042 for Stern (2.28 cites per citing article vs. 1.58). We can also note that 11% of articles citing Öztürk are from him and co-authors, whereas only 0.6% for Stern. One finds similar patterns when extending the analysis to the most cited authors: Shahbaz (2.28 cites per citing article, 1.18 cites per citing author) is very close to Öztürk, whereas Dinda is close to Stern (1.13 cites per citing article, 0.56 cites per citing author).

⁹Information available at <https://www.cag.edu.tr/en/academic-staff/104/about>

All this points to distinct practices in scientific writing and publishing. Statistics of table 4 confirm this trend on a broader scale. Block A presents the highest number of articles per author, while B and C have lower similar values at close to half that of A. Yet, blocks A and B correspond to authors with a longer activity span and who started publishing earlier in the field. The average number of yearly articles shows that, on one side, blocks A and B (around Stern) are close to twice less prolific than block C (D representing again a middle ground). Just as we observed on Stern vs. Öztürk, blocks A and B have a lower number of unique co-authors than block C and, noticeably, D. In a nutshell, the block around Öztürk exhibits distinct publication behaviors: its authors arrived later, yet they publish more often, which contributes to inflating the number of articles, and with more distinct people.

3.3 A shift in journals

We observe a similar divergence when looking at publication outlets. Overall, we found 733 unique outlets in our corpus, with a mean of 3.7 articles per outlet. As is common, we have a large dispersion and a long tail of 642 outlets with no more than 4 articles, representing 35% of the corpus. On the contrary, there are 18 outlets that have published at least 20 articles on the EKC, and these 18 journals represent together 42% of our corpus. We list these journals in table 6, along with their temporal profiles which describe for each journal the proportion of EKC articles that fall in each of the three time periods.

These temporal profiles deliver an interesting picture. The debate on EKC has originated within studies of development as evidenced by the early importance of a journal such as *Environment and Development economics*. It has further taken roots in environmental or energy economics with journals including *Environment and Resources Economics*, *Ecological Economics* or *Energy Policy*. If these journals still publish articles on EKC, the literature from the 2010s is more and more published in journals that are engineering or natural sciences oriented, such as *Environmental Science and Pollution Research* or *Science of the Total Environment*.

These temporal profiles signal a shift in journals. The move away from *Ecological economics* is especially clear: 80% of this journal's articles on EKC are published in the first period, before 2011, whereas the field has experienced a tremendous growth afterwards. This is also surprising given the fact that the EKC was identified by a citation analysis of Ma and Stern (2006) as a prominent topic of ecological economics (as represented by the journal *Ecological economics*), as opposed to environmental economics (as represented by the *Journal of Environmental Economics and Management*). In the period 2006-2010, 46 articles of our corpus were indeed published in *Ecological economics*, but the flow of publications strongly diminishes afterwards, and the output of EKC research was directed to other journals, and actually outside the sphere of economic research. We can see that move if we classify journals in Table 6 as economic journals, when the matter classification of the journal in Scopus contains the word "economics", and non-economic journals otherwise¹⁰. For the articles published

¹⁰For the journals of Table 6, the journals classified as economic are these with 'economic*' in the title, except for *Environ. Dev. Sustainability*, which is also classified as economic.

	<i>Journal</i>	#A	Temporal profile	% Articles citing Stern	Öztürk	<i>r</i>
1989	<i>Ecol. Econ.</i>	118		0.63	0.02	0.97
1991	<i>Environ. Resour. Econ.</i>	38		0.50	0.03	0.95
1996	<i>Environ. Dev. Econ.</i>	36		0.56	0.03	0.95
2001	<i>Int. J. Global Environ. Iss.</i>	24		0.42	0.04	0.91
1984	<i>Economic Modelling</i>	20		0.50	0.05	0.91
1973	<i>Energy Policy</i>	75		0.64	0.28	0.70
1979	<i>Energy Economics</i>	52		0.50	0.31	0.62
1999	<i>Environ. Dev. Sustainability</i>	28		0.54	0.39	0.58
1970	<i>J. Environmental Management</i>	20		0.40	0.30	0.57
1997	<i>Renew. Sustainable Energy Rev</i>	61		0.61	0.56	0.52
2009	<i>Sustainability</i>	78		0.35	0.40	0.47
1976	<i>Energy</i>	30		0.40	0.50	0.44
2001	<i>Ecological Indicators</i>	49		0.41	0.57	0.42
2008	<i>Energies</i>	23		0.35	0.52	0.40
1993	<i>J. Cleaner Production</i>	102		0.30	0.49	0.38
2011	<i>Int. J. Energy Econ. Policy</i>	57		0.39	0.77	0.33
1994	<i>Environ. Sci. Pollution Res.</i>	286		0.30	0.72	0.30
1972	<i>Sci. Total Environment</i>	52		0.17	0.50	0.26

Table 6: Journals publishing at least 20 articles, preceded by year of foundation (first volume). #A indicates the total number of articles in the corpus. The proportion of articles citing respectively Stern and Öztürk is shown, as well as the ratio r of Stern to Stern + Öztürk citations, which is used to rank the table. Temporal profiles are bar charts of the number of articles published over the three periods (1995-2011, 2012-16, 2017-21). All profiles are scaled to their maximum value.

in journals of Table 6, 72% were published in economic journals over the first period 1995-2011, 26% over the period 2012-2016 and only 11% over the last period 2017-2021. The move towards non-economic journal is even more pronounced when one zooms in the first period: 96% of the articles published in journals of Table 6 were in economic journals over the period 1996-2000, 84% over the period 2001-2005, and 63% over the period 2006-2010. The EKC research thus changes significantly over the whole period and was found more and more relevant for fields other than economics. The growth of EKC articles that we have observed is actually attributable to journals classified as non-economic. The field has actually expanded outside economics.




Figure 5: Relative number of articles within Stern’s and Öztürk’s worlds of references. The chart starts in 1999 as there are very few articles before that year and they are all published in Stern’s world.

To get a better sense of this evolution, we go back to Stern and Öztürk. Table 6 indicates the proportion of articles citing either Stern and Öztürk for each journal, as well as the ratio $r \in [0, 1]$ between Stern and Stern+Öztürk proportions — a ratio close to 1 indicates a dominance of articles citing Stern. The table is actually ordered by descending value of r . Articles in journals on top, mostly environmental economics journals, cite far more Stern than Öztürk, whereas articles on the bottom cite far more Öztürk. The shape of temporal profiles exhibits a strong correlation with r in the sense that higher r ratios visibly correspond to earlier periods of EKC publications.

Furthermore, we define the *world of references* of Stern and of Öztürk: If Stern cites an article that is published in journal J , then J is in the world of references of Stern, and similarly for Öztürk. This aims at grasping the journals that an author is aware of and that he considers worth citing. Simply examining the sets gives interesting results because for example some journals like *J. Cleaner Production* or *Renewable Sustainable Energy Rev* are absent from Stern’s world of references. The temporal count in figure 5 is even more striking. We compute the yearly distribution of the ratio of articles published in each world of references relative to the total number of articles published that year in our corpus. We can map the sliding of the field as more and more articles are published in Öztürk’s world of references and less and less in Stern’s one. This reflects the growing contribution in the last time period of journals such as *Environ. Sci. Pollution Res.* or *Sustainability* (from the controversial publisher MDPI) that are principally active in the last period yet are also not in Stern’s world of references. Finally, this imbalance does not concern only the two authors under scrutiny but also extends to journal publication patterns observed in the broader blocks.

<i>Block</i>	“Stern journals”	“Öztürk journals”
<i>A</i>	.571	.111
<i>B</i>	.161	.164
<i>C</i>	.007	.564
<i>D</i>	.022	.436

Table 7: Ratio of articles published in Öztürk and Stern “world of references” journals for each block.

Table 7 shows the share of articles that are published by authors of each block in the respective Öztürk vs. Stern world of references journals: blocks C and D almost do not publish in Stern journals, and do publish massively in Öztürk journals, whereas block A principally publishes in Stern journals and block B, interestingly, exhibits a balanced diet of publication between Stern and Öztürk journals.

4 Discussion

Our results hint at the coexistence of two main epistemic practices with relatively distinct citation and, accordingly, co-authorship networks, also centered around two distinct periods of time, while a shift exists not only in terms of publication venues but also topics, while the relatively balanced share of positive vs. negative results indicates ongoing debate on the empirical validation of the EKC. This further hints at the co-existence of distinct epistemic communities, which we understand here as a community of actors who share a goal of knowledge creation within a common set of topics, beliefs, rules, and, in fine, practices. This notion of epistemic community builds upon two main strands: in political science, where the concept has been introduced (Ruggie, 1975; Haas, 1992) and in social epistemology (Schmitt, 1995; Giere, 2002), whereby an epistemic community corresponds to a group of scientists who acknowledge a specific collection of conceptual tools and representations, within a “paradigm”, may collaborate on specialized tasks, possibly working in a decentralized manner. The joint notion of a shared subset of knowledge issues and procedural authority that characterizes epistemic communities (Cowan et al., 2000) is particularly relevant in our context, where we focus on a main issue (the EKC), various subtopics, and diverse blocks of actors defined by shared citation or reference patterns (Cohendet et al., 2001).

We would like to offer several discussion points that also relate to some limitations of our approach. For one, we see here that the algorithm tends to under-report negative results. Despite the correction that we apply, this sheds light on a broader issue. In effect, articles questioning the rationale behind the EKC and the methods used to find it are not classified as negative results, even though they could be regarded so as they express reservations on the possibility of existence of the object under investigation.

Some further caution is warranted to interpret the number of positive vs. negative results. The following remarks do not address the link between the sentences in the abstract and what is found by the algorithm, but rather the various links between the

research and the sentences of abstracts.

First, there is a general tendency in research to over-publish significant results vs. null results (e.g., Sterling, 1959; Sterling et al., 1995; Ferraro and Shukla, 2020). Null results here would mean the absence of any meaningful relationship between, say, GDP and a pollutant. In the EKC context, however, the publication bias for positive results does not seem relevant: given how the very existence of an EKC for various pollutants has been framed from the very beginning, null results are likely to be conceived as results that invalidate the EKC hypothesis and, thus, published.

Second, there is a reporting bias in the text itself, namely a potential discrepancy between what is found by the authors in the article and what is reported in the abstract. The abstract may report positive results whereas the article actually finds mixed results. A somehow different example of reporting bias is found in Apergis and Ozturk (2015) (cited 187 times, 9th most cited article). The abstract reports “empirical support to the presence of an Environmental Kuznets Curve hypothesis”. However, the article investigates the EKC thanks to a regression involving up to the third power of GDP. The coefficient of the first power is positive and coefficient of the second power is negative, hence the claim that “the estimates have the expected signs” for an EKC. However, the coefficient of the third power is positive, which shows the presence of an N-shape, usually considered a refutation of an EKC. The article offers no explanations of why the third power has been overlooked.

It is also interesting to contrast the methodology followed and the results it delivers with an assessment from an expert in the field. Given his activity time span and dominance in citations patterns coming from all blocks, as the analysis has shown, David I. Stern may be considered the leading expert. As stated in his recent review, Stern (2017) is quite critical of the EKC hypothesis. He shows raw data on emissions of CO₂ per capita and GDP per capita and concludes “there is little sign of an EKC effect” (Stern, 2017, 18). Following the argument already developed in Stern (2010), he estimates that “the relationship between the levels of [...] CO₂ emissions and income per capita is monotonic when the effect of the passage of time is controlled for” (Stern, 2017, 15). A decreasing trend is thus attributed to a time-effect. This skepticism towards EKC reflects a long-standing position. Indeed, the first article (Stern et al., 1996, which was not concerned by CO₂) already emphasizes the mixed results of the empirical investigation of the EKC hypothesis, as well as the estimation problems.

In this regard, the longest-lasting leading expert in the field would appear to support a view that goes against numerous positive results from the recent literature. There could be several explanations to explain this discrepancy. One would be that the expert holds a minority position in the field, but another one would be that there are, in fact, two fields. As we have seen in the previous section, there have been changes in the recent literature. The increase in articles on EKC in the last decade comes from a different breed of journals, with a topical focus that appears to be somewhat distinct from the earlier wave of journals.

Furthermore, our query on Scopus makes the implicit assumption that the term “environmental Kuznets curve” is well-defined in the literature, so that all articles employing it actually speak of the same object. The term seems to be quite narrow, so that we can have at first sight the impression that the assumption is truly warranted. However, a more detailed analysis shows that this is actually questionable. At the very

beginning, the EKC was envisioned as a bivariate relationship between GDP and environmental pressure (depending on pollutant). This is what is reflected in most of the definitions, as the one we provided, that insist on the relationship between GDP on the one hand and environmental degradation on the other. It is also this framing that has been adopted constantly by Stern (Stern et al., 1996; Stern, 2004), more recently see the definition in Stern (2017, p. 8) and the standard EKC regression model (Stern, 2017, p. 13).

Notwithstanding, there have been considerable changes over the three decades that span the scientific existence of the subject. An examination of the articles that found an EKC for CO₂ shows that they rarely consider a bivariate relationship between CO₂ and GDP. A large number of articles are more generally concerned with what they called the emissions / energy and growth nexus. That is, they considered the multivariate relationship between energy, emissions and GDP at least. For example Apergis and Payne (2009) consider the causal relationship between CO₂ emissions, energy consumption, and output. Within this relationship, the relationship between carbon dioxide and GDP, with energy consumption controlled for, is still called an EKC when it is of an inverted U shape. The validity of the EKC is thus tested within a setting different from the one framed by Stern, and the meaning of the EKC has changed.

More generally, further control variables have been introduced, economic or political (institutions), which singularly complicates the meaning and the interpretation of the EKC. New statistical methods, new datasets, new control variables, new environmental stressors, all have stretched the object EKC in different directions, to the point of strongly diminishing its consistency.

The multiple methods available to assess the existence of the EKC also add to the confusion. For example, for the already discussed study on EKC for CO₂ in Vietnam, Al-Mulali et al. (2015) did not confirm the EKC hypothesis, because the elasticity of CO₂ emission with respect to GDP is positive both the short and long run. They rely on the framework set out by Narayan and Narayan (2010) who have however a different understanding of what constitutes a validation of the EKC. For them, the EKC is validated as long as the long-run elasticity is lower than the short-run elasticity. One can easily fall into traps when interpreting this kind of results.

In this regard, the blocks and waves that we exhibited hint at the possibility that what we called the field of the EKC research may actually be subdivided into subfields that share the use of the term “EKC” and references to seminal articles, yet may diverge in what is meant by it. Thus, our findings on the position occupied by Stern put his assessment into perspective.

Conversely, Stern’s assessment of the field and the justifications he put forward points to some limitations of the methods employed here. When counting positive and negative results, we implicitly accept that the underlying articles have used a sound and meaningful method to reach their results, that is that the positive or negative claims made in the abstract are actually genuine results. David I. Stern has however constantly emphasized econometric problems in estimating the EKC. Stern and Common (2001) find that results of regressions are likely spurious and that misspecifications are pervasive. In 2004, he concluded that “the EKC results have a very flimsy statistical foundation” (Stern, 2004). According to him, the field is still plagued with “the naïve econometric approaches used in much of the literature” (Stern, 2017, p. 24). This view

is not unique to David I. Stern and can also be found, for example, in Carson (2010) the unique article on the EKC that was published in the *Review of Environmental Economics and Policy*, an official journal of the American and European associations of resource and environmental economics.

This reminds that our semantical analysis detects claims, and that claims and results cannot be equated without a critical evaluation of the methodology that substantiates these claims, an evaluation that can usually only be provided by field experts. Indeed, the critical assessment by Stern goes beyond the impression we can get as external observers that some results may not be reliable – for instance by employing flawed econometric tools as was pointed by Wagner (2008, 2015). An external assessment has to factor in this reliability, albeit this is very difficult to do without being an expert of the techniques employed in the literature. In our perspective, it is outside of the scope of this article to develop systematic rules to appraise the methodological aspects of articles on an automated basis.

In a way, the evolution that we observed on Figure 2 essentially comes down to a quantitative confirmation that there is still an active debate between positive and negative claims around phenomena that are said, by authors, to be connected to the concept of the environmental Kuznets curve. We leave to further research the possibility of characterizing in more detail the approaches and techniques that are typically in use in each of the research streams that we identified, and the related semantic shifts on “EKC”.

5 Conclusion

Our analysis of the field structured around the use of “environmental Kuznets curve” combines two very recent computational methods in an integrated fashion: semantic and network analyses. For the former, the use of semantic hypergraphs enables us to go beyond lexicometric patterns and to extract both positive and negative claims about the validity of the EKC hypothesis from abstracts. For the latter, using degree-corrected stochastic block-modeling reveals the structure of the author citation network as a compact meta-graph made of a few blocks and easily-interpretable connections between them (i.e., meta-graph nodes are blocks of authors and meta-graph edges are sets of citations between members of blocks).

The combination of topological and semantic features, and a variety of other metrics, both temporal and structural, converges on a characterization of the field that reveals, in essence, the existence of two epistemic communities: one roughly centered around Stern, a long-lasting expert of the field, and one around Öztürk, a more recent expert that also currently dominates the field in terms of citation counts. There also appears to be a remarkable temporal and, to a lesser extent, topical discontinuity between these communities. The first wave and epistemic community, centered around the Stern block, is on the whole less positive on EKC, publishes less often and is less endogamic and is more focused on SO_x and NO_x. The second wave and epistemic community, consisting of the Öztürk block, publishes more positive results and more results overall, is more focused on GHG and energy and is more endogamic. There is also a smaller epistemic community dominated by Chinese authors and focused on

China, that appears to be a middle ground between the two waves according to the various metrics. The divergence of the two communities is also apparent in publication venues — what we call “worlds of references” that are quite distinct and whose activity closely follows the two temporal waves.

We can risk an explanation of the drift between two communities and the shift in main journals we have observed. The research on the EKC that originates in the development and environmental economics community was very intense in the first years, with lively debates about the existence of the EKC and its policy implications. However, after more than fifteen years of research, part of the economic community became suspicious about the EKC, as evidenced for example by the article published by Carson (2010) in the official journal of the American and European associations of resource and environmental economics, that precisely aims at synthesizing lessons for policy-making learned from research of this community. In those circles, EKC research was no longer fashionable as it had been. Newcomers that were attracted by this research, with a different focus on energy and GHG emissions, appear to have somewhat migrated to more engineering-oriented journals, at the crossroads of environmental sciences and energy research. This is the place where the literature currently experiences a massive increase. Whether this increase will continue at the outskirts of economics or whether it foretells a triumphal comeback of the EKC into economics is an open question. The case of the EKC research tells us that research topics have a life of their own, across disciplinary boundaries, that they can survive and even thrive, even when the passions that animated the original players of the field have faded away for scientific reasons.

If we now focus on claims about the EKC, we observe on the whole that the share of positive claims as reported in abstracts has consistently increased over the years, yet remains of the same order of magnitude as negative claims — the debate is not closed. Furthermore, our appraisal of positive vs. negative claims is based on authors’ reporting in abstracts: casual examination of article contents reveals that some positive abstracts correspond to more nuanced results in the article itself, whereby negative findings are intertwined with positive ones, as well as nuances of what counts as an EKC and which variables should be considered (for instance in terms of economic development, pollutants, and sets of countries). Beyond distinct publication and citation practices, this might more broadly suggest that the discontinuity and difference in results may be related to different understandings of the EKC, and the scientific areas, methodologies and topics that are relevant to its appraisal and validation.

Acknowledgments. This research has been partially supported by the “Socseemics” Consolidator grant funded by the European Research Council (ERC) under the European Union Horizon 2020 research and innovation program, grant agreement No. 772743.

References

Acaravci, Ali and Ilhan Ozturk. 2010. On the relationship between energy consumption, CO₂ emissions and economic growth in Europe. *Energy*, 35(12): 5412–5420.

- Ahmed, Nesreen, Jennifer Neville, and Ramana Rao Kompella. 2011. Network Sampling via Edge-based Node Selection with Graph Induction. Purdue University.
- Al-Mulali, Usama, Behnaz Saboori, and Ilhan Ozturk. 2015. Investigating the environmental Kuznets curve hypothesis in Vietnam. *Energy Policy*, 76: 123–131.
- Andreoni, James and Arik Levinson. 2001. The simple analytics of the environmental Kuznets curve. *Journal of Public Economics*, 80(2): 269–286.
- Ang, James B. 2007. CO2 emissions, energy consumption, and output in France. *Energy Policy*, 35(10): 4772–4778.
- Anwar, Muhammad Azfar, Qingyu Zhang, Fahad Asmi, Nazim Hussain, Auke Plantinga, Muhammad Wasif Zafar, and Avik Sinha. 2022. Global perspectives on environmental kuznets curve: A bibliometric review. *Gondwana Research*, 103: 135–145.
- Apergis, Nicholas and Ilhan Ozturk. 2015. Testing environmental Kuznets curve hypothesis in Asian countries. *Ecological indicators*, 52: 16–22.
- Apergis, Nicholas and James E. Payne. 2009. CO2 emissions, energy usage, and output in Central America. *Energy Policy*, 37(8): 3282–3286.
- Barbier, Edward B. 1997. Introduction to the environmental Kuznets curve special issue. *Environment and Development Economics*, 2(4): 369–81.
- Bashir, Muhammad Farhan, Benjiang Ma, Muhammad Adnan Bashir, and Luqman Shahzad. 2021. Scientific data-driven evaluation of academic publications on environmental Kuznets curve. *Environmental Science and Pollution Research*, 28: 16982–16999.
- Beckerman, Wilfred. 1992. Economic growth and the environment: Whose growth? whose environment? *World Development*, 20(4): 481–496.
- Bonaccorsi, Andrea, Nicola Melluso, Filippo Chiarello, and Gualtiero Fantoni. 2021. The credibility of research impact statements: A new analysis of REF with Semantic Hypergraphs. *Science and Public Policy*, 48(2): 212–225.
- Brock, William A. and M. Scott Taylor. 2010. The Green Solow model. *Journal of Economic Growth*, 15(2): 127–153.
- Brundtland, Gro Harlem, Mansour Khalid, Susanna Agnelli, Sali Al-Athel, Bernard Chidzero, Lamina Fadika, Volker Hauff, Istvan Lang, Ma Shijun, Margarita Morino de Botero et al. 1987. *Our common future*. Oxford University Press.
- Carson, Richard T. 2010. The Environmental Kuznets Curve: Seeking Empirical Regularity and Theoretical Structure. *Review of Environmental Economics and Policy*, 4(1): 3–23.
- Claveau, François and Yves Gingras. 2016. Macrodynamics of Economics: A Bibliometric History. *History of Political Economy*, 48(4): 551–592.

- Claveau, François and Catherine Herfeld. 2018. Network Analysis in the History of Economics. *History of Political Economy*, 50(3): 597–603.
- Cohendet, Patrick, Frédéric Créplet, and Olivier Dupouet. 2001. *Economics with Heterogeneous Interacting Agents*, Berlin: Springer, chapter Organisational innovation, communities of practice and epistemic communities: the case of Linux. 303–326.
- Cowan, Robin, Paul A. David, and Dominique Foray. 2000. The Explicit Economics of Knowledge Codification and Tacitness. *Industrial & Corporate Change*, 9(2): 212–253.
- Cropper, Maureen and Charles Griffiths. 1994. The interaction of population growth and environmental quality. *The American Economic Review*, 84(2): 250–254.
- Dinda, Soumyananda. 2004. Environmental Kuznets Curve Hypothesis: A Survey. *Ecological Economics*, 49(4): 431–455.
- Ettinger, Allyson. 2020. What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics*, 8: 34–48.
- Ferraro, Paul J. and Pallavi Shukla. 2020. Is a Replicability Crisis on the Horizon for Environmental and Resource Economics? *Review of Environmental Economics and Policy*, 14(2): 339–351.
- Giere, R. 2002. Scientific Cognition as Distributed Cognition. In Peter Carruthers, Stephen Stich, and Michael Siegal (eds.), *The Cognitive Basis of Science*, Cambridge University Press. 285–299.
- Grossman, Gene M. and Alan B. Krueger. 1991. Environmental Impacts of a North American Free Trade Agreement. Cambridge (Ma.): NBER. Working Paper 3914.
- Haas, P. 1992. Introduction: epistemic communities and international policy coordination. *International Organization*, 46(1): 1–35.
- Haberl, Helmut, Dominik Wiedenhofer, Doris Virág, Gerald Kalt, Barbara Plank, Paul Brockway, Tomer Fishman, Daniel Hausknost, Fridolin Krausmann, Bartholomäus Leon-Gruchalski, Andreas Mayer, Melanie Pichler, Anke Schaffartzik, Tânia Sousa, Jan Streeck, and Felix Creutzig. 2020. A systematic review of the evidence on decoupling of GDP, resource use and GHG emissions, part II: synthesizing the insights. *Environmental Research Letters*, 15(6): 065003.
- Herfeld, Catherine and Doehne Malte. 2018. Five reasons for the use of network analysis in the history of economics. *Journal of Economic Methodology*, 25(4): 311–328.
- Holland, Paul W, Kathryn Blackmond Laskey, and Samuel Leinhardt. 1983. Stochastic blockmodels: First steps. *Social networks*, 5(2): 109–137.
- Jalali, Zeinab S, Alireza Rezvanian, and Mohammad Reza Meybodi. 2016. Social network sampling using spanning trees. *International Journal of Modern Physics C*, 27(05): 1650052.

- Jia, Yuntao, Jared Hoberock, Michael Garland, and John Hart. 2008. On the visualization of social and other scale-free networks. *IEEE transactions on visualization and computer graphics*, 14(6): 1285–1292.
- Karrer, Brian and Mark EJ Newman. 2011. Stochastic blockmodels and community structure in networks. *Physical review E*, 83(1): 016107.
- Kijima, Masaaki, Katsumasa Nishide, and Atsuyuki Ohyama. 2010. Economic models for the environmental Kuznets curve: A survey. *Journal of Economic Dynamics and Control*, 34(7): 1187–1201.
- Koondhar, Mansoor Ahmed, Muhammad Shahbaz, Kamran Ali Memon, Ilhan Ozturk, and Rong Kong. 2021. A visualization review analysis of the last two decades for environmental Kuznets curve “EKC” based on co-citation analysis theory and pathfinder network scaling algorithms. *Environmental Science and Pollution Research*, 28: 16690–16706.
- Kuznets, Simon. 1955. Economic Growth and Income Inequality. *The American Economic Review*, 45(1): 1–28.
- Leskovec, Jure and Christos Faloutsos. 2006. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. 631–636.
- Ma, Chunbo and David I. Stern. 2006. Environmental and ecological economics: A citation analysis. *Ecological Economics*, 58(3): 491–506.
- Maiya, Arun S and Tanya Y Berger-Wolf. 2011. Benefits of bias: Towards better characterization of network sampling. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. 105–113.
- Meadows, Donella H., Dennis L. Meadows, Jørgen Randers, and William W. Behrens. 1972. *The limits to growth: a report for the Club of Rome’s project on the predicament of mankind*. New-York: Universe Books.
- Menezes, Telmo and Camille Roth. 2019. Semantic Hypergraphs. *arXiv*, 1908.10784.
- Narayan, Paresh Kumar and Seema Narayan. 2010. Carbon dioxide emissions and economic growth: Panel data evidence from developing countries. *Energy Policy*, 38(1): 661–666.
- Naveed, Amjad, Nisar Ahmad, Reza FathollahZadeh Aghdam, and Angeliki N Menegaki. 2022. What have we learned from Environmental Kuznets Curve hypothesis? A citation-based systematic literature review and content analysis. *Energy Strategy Reviews*, 44: 100946.
- Nordhaus, William D. 1992. Lethal Model 2: The Limits to Growth Revisited. *Brookings Papers on Economic Activity*, 1992(2): 1–59.

- Panayotou, Theodore. 1993. Empirical Tests and Policy Analysis of Environmental Degradation at Different Stages of Economic Development. Geneva: International Labour Office.
- Panayotou, Theodore. 1997. Demystifying the environmental Kuznets curve: turning a black box into a policy tool. *Environment and Development Economics*, 2(4): 465–484.
- Peixoto, Tiago P. 2018. Nonparametric weighted stochastic block models. *Physical Review E*, 97(1): 012306.
- Piketty, Thomas. 2014. *Capital in the twenty-first century*. Cambridge Massachusetts: The Belknap Press of Harvard University Press.
- Ribeiro, Bruno and Don Towsley. 2010. Estimating and sampling graphs with multidimensional random walks. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. 390–403.
- Ruggie, John Gerard. 1975. International Responses to Technology: Concepts and Trends. *International Organization*, 29(3): 557–583.
- Sarkodie, Samuel Asumadu and Vladimir Strezov. 2019. A review on Environmental Kuznets Curve hypothesis using bibliometric and meta-analysis. *Science of The Total Environment*, 649: 128–145.
- Schmid, Moritz S, Robert K Cowen, Kelly Robinson, Jessica Y Luo, Christian Briseño-Avena, and Su Sponaugle. 2020. Prey and predator overlap at the edge of a mesoscale eddy: Fine-scale, in-situ distributions to inform our understanding of oceanographic processes. *Scientific Reports*, 10(1): 921.
- Schmitt, Frederik (ed.). 1995. *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham, MD: Rowman & Littlefield.
- Shafik, Nemat and Sushenjit Bandyopadhyay. 1992. Economic Growth and Environmental Quality: Time Series and Crosscountry Evidence. Background Paper for the World Development Report 1992. Washington: The World Bank.
- Shahbaz, Muhammad and Avik Sinha. 2019. Environmental Kuznets curve for CO2 emissions: a literature survey. *Journal of Economic Studies*, 46(1): 106–168.
- Sterling, T. D., W. L. Rosenbaum, and J. J. Weinkam. 1995. Publication Decisions Revisited: The Effect of the Outcome of Statistical Tests on the Decision to Publish and Vice Versa. *The American Statistician*, 49(1): 108–112.
- Sterling, Theodore D. 1959. Publication Decisions and their Possible Effects on Inferences Drawn from Tests of Significance—or Vice Versa. *Journal of the American Statistical Association*, 54(285): 30–34.
- Stern, David I. 2004. The Rise and Fall of the Environmental Kuznets Curve. *World Development*, 32(8): 1419–1439.

- Stern, David I. 2010. Between estimates of the emissions-income elasticity. *Ecological Economics*, 69(11): 2173–2182.
- Stern, David I. 2017. The environmental Kuznets curve after 25 years. *Journal of Bioeconomics*, 19(1): 7–28.
- Stern, David I. and Michael S. Common. 2001. Is There an Environmental Kuznets Curve for Sulfur? *Journal of Environmental Economics and Management*, 41(2): 162–178.
- Stern, David I., Michael S. Common, and Edward B. Barbier. 1996. Economic growth and environmental degradation: The environmental Kuznets curve and sustainable development. *World Development*, 24(7): 1151–1160.
- Van Ham, Frank and Martin Wattenberg. 2008. Centrality based visualization of small world graphs. *Computer Graphics Forum*, 27(3): 975–982.
- Voudigari, Elli, Nikos Salamanos, Theodore Papageorgiou, and Emmanuel J Yannakoudakis. 2016. Rank degree: An efficient algorithm for graph sampling. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 120–129.
- Wagner, Martin. 2008. The carbon Kuznets curve: A cloudy picture emitted by bad econometrics? *Resource and Energy Economics*, 30(3): 388–408.
- Wagner, Martin. 2015. The Environmental Kuznets Curve, Cointegration and Nonlinearity. *Journal of Applied Econometrics*, 30(6): 948–967.
- World Bank. 1992. World Development Report 1992. Development and the Environment. New York: World Bank.
- Zhao, Ying, Haojin Jiang, Yaqi Qin, Huixuan Xie, Yitao Wu, Shixia Liu, Zhiguang Zhou, Jiazhi Xia, Fangfang Zhou et al. 2020. Preserving minority structures in graph sampling. *IEEE Transactions on Visualization and Computer Graphics*, 27(2): 1698–1708.

A Semantic hypergraphs

We provide here a brief theoretical and practical overview of SH, especially in the context of this article. Interested readers may find further details in (Menezes and Roth, 2019).

First, SH are made of (semantic) hyperedges, representing utterances as relations consisting of a predicate (annotated with /P) followed by an arbitrary number of participants (concepts are annotated with /C). The building blocks of SH are *atoms*, which mostly correspond to words annotated with types (a few additional atoms are defined, for example to define connectors for compound nouns). The sentence “Mary plays chess.” translates to:

(plays/P mary/C chess/C)

Relations can be nested, for example in “John says Mary plays chess.”:

(says/P john/C (plays/P mary/C chess/C))

Relations can be combined with conjunctions (annotated with /J), for example in “John reads and Mary plays chess.”:

(and/J (reads/P john/C) (plays/P mary/C chess/C))

Relations connect concepts, but concepts themselves can be recursively constructed with other concepts with the help of builders (/B), such as in “Mayor of Berlin”:

(of/B mayor/C berlin/C)

Finally, concepts, predicates and other hyperedges can be made more specific with modifiers (/M):

(first/M (of/B mayor/C berlin/C))

The above recursivity is an essential property of SHs in that they mirror the recursive scaffolding of natural language sentences:

(plays/P (the/M (last/M (of/B mayor/C berlin/C))) (chess/C))

Defining hypergraphic rules to detect EKC-related results. We used Graphbrain¹¹, an open source library that implements a variety of functions related to SH, to convert all the sentences in article abstracts to hyperedges. Then, we took advantage of the structure and regularity imposed on natural language by the SH representation to define rules that detect the reporting of results either confirming or refuting EKCs.

We focus on relations with predicates that are likely to be related with reporting results, e.g.: “found”.

$$H = (p \ s \ a \ a' \ \dots) \quad (1)$$

where p is of type predicate (/P), and the (optional) arguments are denoted as s for the potential subject of the relation, and a , a' , etc. for non-subject arguments. Graphbrain is particularly efficient at determining the semantic roles fulfilled by the sequence of arguments of a relation hyperedge (we refer the interested reader to Menezes and Roth, 2019, for more details).

We then iteratively refined a small set of rules to extract hyperedges denoting results related to the empirical confirmation or refutation of EKC. To this end, we used a subset of the corpus made of 500 randomly selected sentences, out of 23042 (2.17%), akin to a training set in conventional machine learning. The main steps were as follows:

- We looked for the most commonly used predicates and extracted the ones usually associated with results: we explored 100 top predicates using the Graphbrain tool, which offers this functionality, and found 20 of them (shown in table 8) to more frequently correspond to claims about results.

The first main rule is that p contains an atom belonging to this set of so-called claim predicates, either positive (P^+) or negative (P^-).

¹¹<https://graphbrain.net>

Set	Atoms
Claim predicates	show/P, indicate/P, confirm/P, support/P, suggest/P, reveal/P, provide/P, validate/P, exist/P, positive (P^+) demonstrate/P, verify/P, imply/P, illustrate/P, find/P, point/P, exhibit/P, establish/P, obtain/P, hold/P, follow/P
	negative (P^-) reject/P, challenge/P, fail/P
EKC concepts	kuznets/C, ekc/C, turning/C
Curve concepts	curve/C, shape/C, shaped/C
Negative modifiers	not/M, n't/M, no/M, little/M, poor/M
Result concepts (R)	result/C, finding/C, test/C, evidence/C, support/C

Table 8: The various atom sets used in claim detection and classification rules.

- Many results are not directly about EKC, so we limited the rules to explicit references to EKC. We also included rules to capture the notions of “U-curves” and “N-curves” which, in the corpus, generally refer implicitly to EKC.

In practice, we thus defined a helper rule $\mathcal{E}$ as follows: $\mathcal{E}(H)$ true if $\exists a \in H$ such that,

- either $a \in \{\text{kuznets/C, ekc/C, turning/C}\}$
- or $\exists c \in \{\text{curve/C, shape/C, shaped/C}\}$,
such that $a = (u/C \ c)$ or $a = (n/C \ c)$

- However, if the reference to EKC is in the subject, it is not a result (“EKCs reveal (...)” vs. “We reveal the presence of EKC”), and we include a rule where $\mathcal{E}(s)$ must be false.

The following expression formalizes the above rules for identifying a result claim about EKCs, as a logical expression (whereby ‘ $\wedge$ ’ and ‘ $\neg$ ’ denote respectively the ‘AND’ and ‘NOT’ logical operators):

$$\mathcal{C}(H) = (p \cap (P^+ \cup P^-) \neq \emptyset) \wedge (\mathcal{E}(H) \wedge \neg \mathcal{E}((s))) \quad (2)$$

Positive vs. negative results. To determine if a result claim is positive, negative or unknown in regard to EKC validity, we further take advantage of SH structure. We add a helper rule $\mathcal{N}$ to detect negative claims i.e., cases where the predicate is negated (e.g.: “Could not find evidence of...”), or the concept referring to the result is negated (e.g.: “We found no evidence of...”), or the concept referring to EKC is negated (e.g.: “No EKC was found...”). Formally, $\mathcal{N}(H)$ is true if H contains:

- either, both:
 - a negation in its predicate p :
 $p \cap \{\text{not/M, n't/M, no/M, little/M, poor/M}\} \neq \emptyset$

- and a claim predicate or a result concept in one of its arguments:
 $\exists a \in H, a \cap (P \cup R) \neq \emptyset$ or $\mathcal{E}(a)$
 which further relies on a set R of atoms of so-called “result concepts”:

$$R = \{\text{result/C, finding/C, test/C, evidence/C, support/C}\}$$

This set has been constructed by examining the most frequent atoms directly connected to negations (among the top 100 such atoms), similarly to how P has been built.

- or, recursively, an element of H (possibly a hyperedge itself) for which N is true: $\exists h \in H, \mathcal{N}(h)$ is true.

Finally, table 9 shows how a result claim is classified. Notice that we consider a positive finding of an N-curve (i.e., $\exists a \in H, a = (n/C c)$), as a refutation of EKC. We do not consider double negations, as we found them to increase complexity while having no impact on accuracy.

predicate p based on	$\mathcal{N}(H)$	H contains an N-curve	<i>EKC validation claim type</i>
P^+	True	True	X
		False	<i>negative result</i>
	False	True	<i>negative result</i>
		False	<i>positive result</i>
P^-	True		X
		True	X
	False	False	<i>negative result</i>

Table 9: Classification of a claim H as a positive or negative result, or unknown (X) based on the elements of H .