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

Telmo Menezes*[1], Antonin Pottier†[1,2], and Camille Roth‡[1,3]

[1]Centre Marc Bloch Berlin (CNRS/HU), Friedrichstr. 191, 10117 Berlin, Germany
[2]CIRED (Centre international de recherche sur l'environnement et le développement CNRS/EHESS), 54 Bd Raspail, 75007 Paris, France
[3]CAMS (Centre Analyse et Mathématique Sociales, UMR 8557 CNRS/EHESS), 54 Bd Raspail, 75007 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 do here a review of more than 2000 papers that have been published on the EKC. 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 as well as 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.

1 Introduction

Relationship between environmental impacts and economic income is a hotly debated topic in environmental economics. In the 1990s, emerged the idea that environmental impacts of economic activity increase then decrease. The promise behind it was, contrary to the views filled of tensions between economy and the environment that were at play in the debates of sixties and seventies, that the economic development will somehow take care of its own drawbacks, so that the state of the environment, after an initial period of degradation, will automatically improve.

It can be traced back to Grossman and Krueger (1991), who study the environmental impacts of the free-trade agreement between the USA, Canada and Mexico. They

*menezes@cmb.hu-berlin.de
†pottier@centre-cired.fr
‡roth@cmb.hu-berlin.de

observe that the concentrations of some pollutants increase at low income but decrease at high income, suggesting that trade liberalization that would foster economic development could also be good for the environment. This idea was further carried out by Shafik and Bandyopadhyay (1992) who studied more systematically the link between economic growth and environmental quality. This research was given wide outreach as it was a background paper for the *World Development Report* 1992 (World Bank, 1992). The report displays several patterns of environmental indicators depending on country income: monotonically decreasing, monotonically increasing, or increasing then decreasing.

This inverted U-shaped relationship between the environmental degradation (measured by several indicators of pollutants) and economic development (GDP or income) was named the environmental Kuznets curve (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 observed[1]. Many empirical contributions immediately investigated whether such a relationship exists or not, for what type of pollutants, while others pinpoint the problems in estimating its parameters (Stern et al., 1996). The empirical tests were based on cross-section of countries or panel data. Water and (local) air pollution were mostly investigated. The mechanisms that would explain such a relationship were also debated, along the lines suggested by Grossman and Krueger (1991), distinguishing the effects of an expanding scale of the economy, of composition change of output, and of the progress of technologies. Theoretical models were proposed (Andreoni and Levinson, 2001; Brock and Taylor, 2010). The political implications of the existence or not of an EKC were also discussed (Panayotou, 1997; Dinda, 2004). A literature field, both theoretical and empirical, thus builds up around "the environmental Kuznets curve".

Thirty years after the seminal contributions more than 2000 articles have been published that contains "Environmental Kuznets curve". The purpose of this article is to quantitatively investigate this literature with tools from bibliometric analysis and semantic analysis. Bibliometric analysis refers to the study of the various networks that can be derived from the oriented relation that a citation creates between two papers. Semantic analysis refers to the automatic study of the meaning of sentences. We apply semantic analysis to the abstracts of the collected papers.

The paper closest to ours is Sarkodie and Strezov (2019), who have done a bibliometric analysis of the EKC. They mainly investigate author's contributions, citations, source and country, as well as topics. They however do not use it to map the evolution of the field. Anwar et al. (2022) have done a bibliometric review, identifying key articles according to various methodologies. However their results are not interpreted, so that it does not allow a deep understanding of the development of the EKC literature. Our paper intends to go beyond these existing articles and to connect this research to the social studies of economics. As advocated by Claveau and Herfeld (2018), it uses network analysis to describe and visualize the relations between authors active within the field of EKC research. 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.

[1]Ironically, the Kuznets curve seems to be no more than an accident. See Piketty (2014, introduction).

The paper is organized as follows. Section 2 presents the methodology of semantic analysis and the first results in term of temporal development of the field, type of pollutants that are investigated, type of claims that are made. Section 3 introduces the analysis of the citation network, 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 the difference between the assessment done here and the assessment of an expert of the field, followed by a summary of our findings in Section 5.

2 Methods and descriptive results

2.1 Corpus building

We build our corpus from the database Scopus. Scopus is one of the main bibliographic databases, released by the scientific publisher Elsevier. It covers content from more than 25,000 active titles and 7,000 publishers.[2] We query Scopus for papers containing the exact phrase "environmental Kuznets curve" in title, abstract or keywords. The search was performed on June 4, 2021 delivering 2709 papers. We collect all possible data that were furnished by Scopus. Here, the analysis will mainly be focused on title, authors, abstract and citations. We screen random references to see whether the collected papers were mis-attributed. Given that the syntagm "Environmental Kuznets curve" is already quite narrow, we find no paper that addresses a different topic, so we do not dwell further to exclude wrongly included papers, 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 papers have been published in the past four years and a half (since 2017 included), although the first papers are from 1994 (only 2% have been published before 2000 included). This exponential growth is striking as Scopus is not known to exhibit a strong recent bias, contrary to Google Scholar. We find that in the last decade the number of paper 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.

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 a paper 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 refering to number of citations come from this and so have to be understood as internal i.e., most-cited papers are papers that are most cited by the papers of the corpus.

2.2 Semantic analysis

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

[2] https://www.elsevier.com/solutions/scopus/how-scopus-works

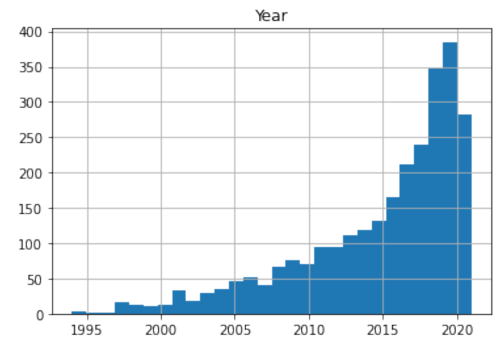

Figure 1: Corpus per year of publication

interesting for an analysis of the field as it allows to reach, to some extent, what an article means, and in turn deepen traditional bibliometric analysis by complementing it with semantic analysis.

Semantic Hypergraphs. 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). Here we just give a short overview of SH. 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 Graphbrain[3], 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 which are most likely to express results. In SH parlance, the (semantic) hyperedge H related to a given relation could be written as:

$$H = (\mathsf{p\ s\ a\ a' \ldots}) \tag{1}$$

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 and found 20 of them (shown in table 1) 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^-).

- 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 $\exists \mathsf{a} \in H$ such that,

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

[3] https://graphbrain.net

Set		Atoms
Claim predicates	*positive (P^+)*	show/P, indicate/P, confirm/P, support/P, suggest/P, reveal/P, provide/P, validate/P, exist/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 1: The various atom sets used in claim detection and classification rules.

- 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 $\mathcal{E}(\mathsf{s})$ must be false.

The following expression formalizes the above rules for identifying a result claim about EKCs:

$$\mathcal{C}(H) = \big(\mathsf{p} \cap (P^+ \cup P^-) \neq \varnothing\big) \wedge (\mathcal{E}(H) \wedge \neg\mathcal{E}((\mathsf{s}))) \tag{2}$$

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, $\mathcal{N}(\mathcal{H})$ is true if H contains:

- either, both:
  - a negation in its predicate p:
    $\mathsf{p} \cap \{\mathsf{not/M}, \mathsf{n't/M}, \mathsf{no/M}, \mathsf{little/M}, \mathsf{poor/M}\} \neq \varnothing$
  - and a claim predicate or a result concept in one of its arguments:
    $\exists \mathsf{a} \in H, \mathsf{a} \cap (P \cup R) \neq \varnothing$ or $\mathcal{E}(\mathsf{a})$
    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 2 shows how a result claim is classified. Notice that we consider a positive finding of an N-curve (i.e., $\exists \mathsf{a} \in H, \mathsf{a} = (\mathsf{n/C\ c})$), 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	*EKC validation claim type*
P^+	True	True	*X*
		False	*negative result*
	False	True	*negative result*
		False	*positive result*
P^-	True		*X*
	False	True	*X*
		False	*negative result*

Table 2: Classification of a claim H as a positive or negative result, or unknown (X) based on the elements of H.

To evaluate the performance of our semantic classifier, we randomly selected 3 sets of 50 abstracts each, to be manually annotated by 3 different persons, so that each person 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 [4] results are shown in table 3.

EKC validation claim type	Precision	Recall
positive result	.809	.847
negative result	.833	.366

Table 3: 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 also, negative results are implied and cannot be found by direct claim

[4] *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.

identification. A telling example is Acaravci and Ozturk (2010), a largely cited paper (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 the validity of environmental Kuznets curve (EKC) hypothesis in Denmark and Italy". However, the paper investigates the relationship between CO_2 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 paper.

The much larger recall error of the algorithm for negative results leads us to correct all values found by the algorithm by multiplying raw quantities of results by the associated precision and dividing them by the associated recall. 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.

We plot the ratio of articles with a positive, a negative, or both a positive and a negative result found in the corpus in figure 2-left. 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 papers in these early stages, as shown in figure 1). As a result, the 1995-2011 period comprises 617 articles; 2012-2016, 621 articles, and 2017-2021, 1467 articles.

We see that positive and negative results are found in similar proportions. Even though there has been a modest albeit increasing lead in the number of papers with positive findings over the recent years, there nonetheless remains a substantial proportion of 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 unconclusive, 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 pressures. Our first task is then to identify which kind of environmental variables have been investigated in a given paper.

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. The pattern can be represented in SH as:

```
(between/B relationship/C (and/J X Y))
```

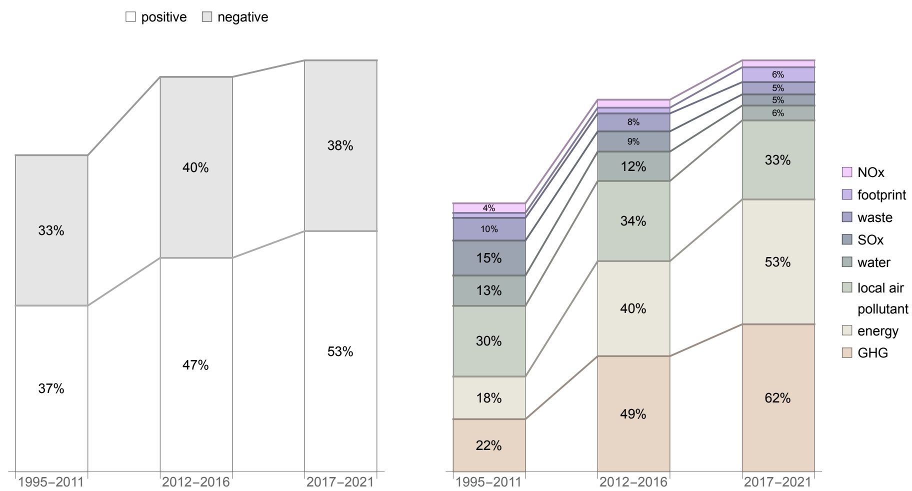

Figure 2: *Left:* estimated percentage of articles with a positive or a negative result (the bars do not add to 100% because articles can present no results, or both positive and negative results). *Right:* percentage of articles evoking a relationship on a given topic (sums over 100% as an article may address several topics, bar elements representing less than 4% were not labeled).

We exploited this pattern to identify all the environmental variables present in X or Y. The most common ones were manually grouped into eight categories: green house gases (labeled as "GHG", which includes for a very large part CO_2-related research), energy, local air pollutants, water, SOx, waste, footprint and NOx.

On the whole, 82% of articles were assigned a topic. On the right of figure 2 we present the ratios of articles mentioning each category. Ratios sum to more than 100% since there is overlap between the categories: a paper 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 there is a large overlap between energy and carbon. This is because the papers 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/CO_2 emissions and energy, which was also found by Sarkodie and Strezov (2019) using a different method (keyword analysis). This is surprising. Although Shafik and Bandyopadhyay (1992) considered CO_2 within their set of pollutants, the early decade of research on EKC was not strongly concerned with CO_2. Dinda (2004)'s review mentions research on CO_2 but focuses on local air pollution and water pollution. On the contrary, from 2014, the papers identified as investigating CO_2 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 local air pollutants such as nitrogen and sulfur oxides (NOx and SOx), and various environmental stressors such as waste, more recent

	Topic							
Period	GHG	energy	local air pollutants	water	SOx	waste	footprint	NOx
1995-2011	1.03	1.29	0.84	0.81	1.18	**1.10**	0.31	0.91
2012-2016	1.05	1.09	1.15	**1.03**	1.00	0.67	0.42	0.92
2017-2021	**1.55**	**1.39**	**1.64**	1.01	**1.39**	1.01	**1.07**	**1.50**
aggregate	*1.35*	*1.31*	*1.32*	*0.95*	*1.19*	*0.94*	*0.88*	*1.14*

Table 4: Evolution of the ratio of papers 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").

publications do not seem to focus explicitly on such specific pollutants: the shares of SOx, NOx, waste as well as water [pollution] strongly decreased; only the share of the generic reference to "local air pollutants" remained stable. Third, the increase of topic-focused research on EKC: the bars are generally increasing which indicates that more and more articles are concerned with relationships involving a specific environmental variable. To summarize, we witness a strong turn to energy and GHG-related pollutants from around 2012.

A cross analysis of positive/negative results vs. topics makes it possible to refine this picture further. Table 4 provides the relative proportion of positive vs. negative results for each topic through time — in other words, it paints the categories of the right side of figure 2 with the categories of its left side. 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, NOx and, to a smaller extent, SOx, experience the strongest increase in positive results. More interestingly, for half of the categories (local air pollutants, footprint, water and NOx), 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. 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 citation network.

To begin, we compute for each author the number of times their papers have been cited. As said before, we only consider citations internal to the corpus. We concentrate on the five most (internally) cited authors, listed in table 5, for whom we can already distinguish several patterns. I. Öztürk and M. Shahbaz entered the fields in the 2010s, and since they have published a large number of papers. D.I. Stern has been present in the field over the all time span, with a large yet lower number of papers per year. Finally, K.-G. Mäler and S. Dinda were active in the 2000s, left the field meanwhile, and have published relatively few papers.

Author	#A	c	Activity time span
Öztürk I.	35	1808	2010-2021
Stern D.I.	17	1644	1996-2020
Shahbaz M.	49	1541	2012-2021
Mäler K.-G.	3	1070	2003-2013
Dinda S.	7	892	2000-2009

Table 5: Five most cited authors. #A is the number of articles, c number of 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. The differences seem to stem from two facts. First Sarkodie and Strezov (2019) use the *Web of Science* database to extract their corpus whereas we rely on *Scopus*. Second, 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.

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. This is a directed network, with each edge indicating that the origin cites the target. The full network is very dense, so to make it tractable 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 an arbitrary in-degree. Furthermore, we consider only authors that received at least 10 citations.

We then use Stochastic Block Models (SBM) (Holland et al., 1983; Peixoto, 2018) and the open-source library *graph-tool* [5] to infer the structure of the network. SBM is a generative model, where each node is assigned to a given block. SBM jointly infers block membership and the matrix of probabilities of connection between nodes in two given blocks (that can be the same), 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 blocks 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

[5] https://graph-tool.skewed.de

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 5 blocks, as depicted in figure 3.

Blocks are ultimately characterized by their linking behavior, which can be interpreted as an indication of their role in the network. For example, blocks A and C both contain a small number of highly cited authors. They do not belong to a single block because there are significant linking behavior differences, both inwards and outwards. A tends to be cited by all blocks and to not cite any other blocks, while C is not cited by one of the big blocks (B) and tends to cite A.

Figure 3(a) shows us that two large blocks (B and D), of comparable size, account for the majority of authors in the network. There is a certain topological symmetry, also in the sense that each one of these blocks surrounds two much smaller ones (A and C), which contain highly cited authors. Notably, Stern belongs to block A and Öztürk to block C. The remaining block (E) is dominated by Chinese authors publishing articles with a strong focus on China. In figure 3(b) 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 appear 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 that enjoys this level of centrality. Another indication of this asymmetry is the more casual observation that Öztürk cites Stern's papers thirteen times, whereas Stern does not cite Öztürk at all.

In table 6 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 networks structure, and therefore help to strengthen the hypothesis that these blocks do indeed correspond to different cultures within EKC research. Blocks A and then B have a low ratio of positive results in comparison to C and D. The ratio of negative results is more similar across blocks, except for A. There is also a clear difference in endogamy, with both C being more endogamic than A and D more than B. Interestingly, block E appears to be a middle ground between A+B and C+D 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 cultures.

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 block C is much more recent and D even more so. This invites the hypothesis of a shift in cultures having taken place during some period in time, possibly around 2012-2015. Finally, the breakdown of the presence of each topic for each block, shown in table 7, also paints a quite heterogeneous picture. It indicates a certain level of focus and specialization proper to some blocks: for instance, A on SOx, C and D on GHG and energy, E on local air pollutants.

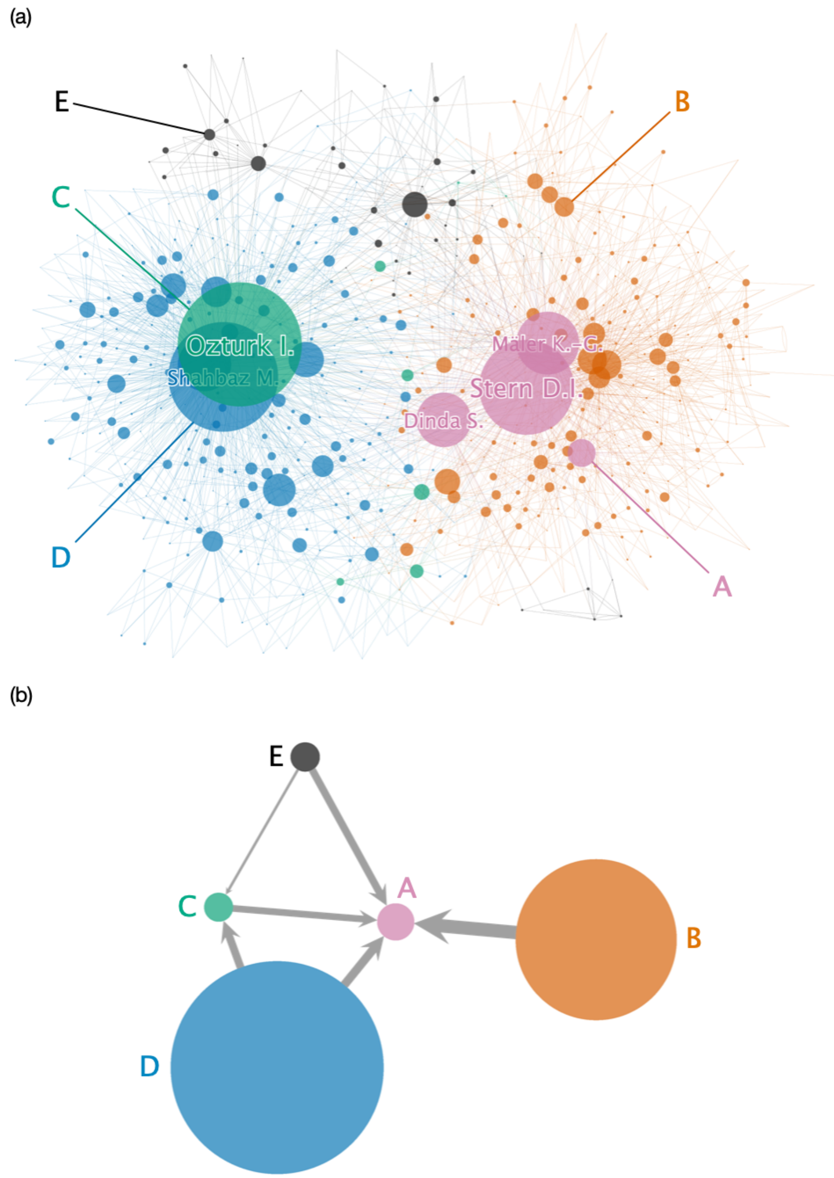

Figure 3: (a) Graph of citations between authors, colored by blocks determined by SBM. Node radius is proportional to in-degree. (b) Graph of author blocks. Edge thickness is proportional to the probability of connection from authors of one block to another. Node radius is proportional to the total number of citations (for the entire corpus) received by the authors in the block.

Block	#authors	pos	neg	endogamy	year	#A	$\overline{\#A}$	k
A	4	.318	.076	.002	2007.7	7.50	0.44	4.5
B	366	.397	.380	.034	2012.4	3.22	0.76	5.1
C	8	.605	.455	.027	2015.9	7.75	1.35	11.1
D	392	.538	.380	.104	2018.1	4.27	1.26	8.9
E	64	.472	.366	.080	2016.6	6.13	1.00	15.6

Table 6: 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).

Block	GHG	energy	local air pollutants	water	SOx	waste	footprint	NOx
A	.300	.100	.233	.000	**.333**	.000	.000	.033
B	.368	.278	.339	.114	.101	.086	.021	**.039**
C	**.717**	**.817**	.183	.033	.000	.017	.050	.000
D	.665	.624	.312	.064	.041	.044	**.092**	.021
E	.418	.352	**.533**	**.146**	.088	**.092**	.015	.038

Table 7: Percentage of articles mentioning each topic for each author block. Bold figures indicate the block where a topic has the highest presence.

3.2 Focusing on the two leaders

David I. Stern, the second most cited author, is one of the most long-lived authors in the field. Indeed, its activity spans 25 years. He began his career[6] 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. He was associate editor of *Ecological Economics* (which appears in Table 8) from 2002 to 2018 and belongs to its editorial board since then. His most cited paper 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 comparatively a rather newer scholar who completed his PhD in Economics in 2009 from Çukurova University in Turke and began publishing on EKC in 2010.[7] He has been

[6] Information available at http://www.sterndavidi.com
[7] Information available at https://www.cag.edu.tr/en/academic-staff/104/about

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 8), founded in 2011 and edited by EconJournals. 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 paper 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 5, they seem to have different scientific practices. As said before, Stern has published less than Öztürk in our corpus, 17 papers against 35, but within 25 years of activity compared to 12, which corresponds respectively to 0.68 papers 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 papers: 1551 distinct authors cite Öztürk vs. 1992 for Stern (hence 1.17 cites per author compared with 0.83), and 794 distinct papers for Öztürk vs. 1042 for Stern (2.28 cites per citing paper vs. 1.58). We can also note that 11% of papers citing Öztürk are from him and co-authors, whereas only 0.6% for Stern. One finds similar patterns when extending the analysis to the five most cited authors: Shahbaz (2.28 cites per citing paper, 1.18 cites per citing author) is very close to Öztürk, whereas Mäler and Dinda are close to Stern (1.05 cites per citing paper, 0.56 cites per citing author and 1.13 cites per citing paper, 0.56 cites per citing author).

All this points to distinct practices in scientific writing and publishing. Statistics of table 6 confirm this trend on a broader scale. Apparently, blocks A and C around leading authors exhibit a similar number of articles per author, about twice higher than blocks B and D, which are also similar in this regard (E being in the middle). 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 papers shows that, on one side, blocks A and B (around Stern) are twice to three times less prolific than blocks C and D (E 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 blocks C, D and, noticeably, E. In a nutshell, blocks around Öztürk exhibit distinct publication behaviors: their authors arrived later, yet they publish more often, which contributes to inflating the number of papers, 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 papers, 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 8, along with their temporal profiles which describe for each journal the proportion of EKC papers that fall in each of the three time periods.

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

Table 8: 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 Okturk is shown, as well as the ratio r of Stern to Stern + Öztürk citations. 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.

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 papers 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*.

To get a better sense of this evolution, we go back to Stern and Öztürk. Table 8 in-

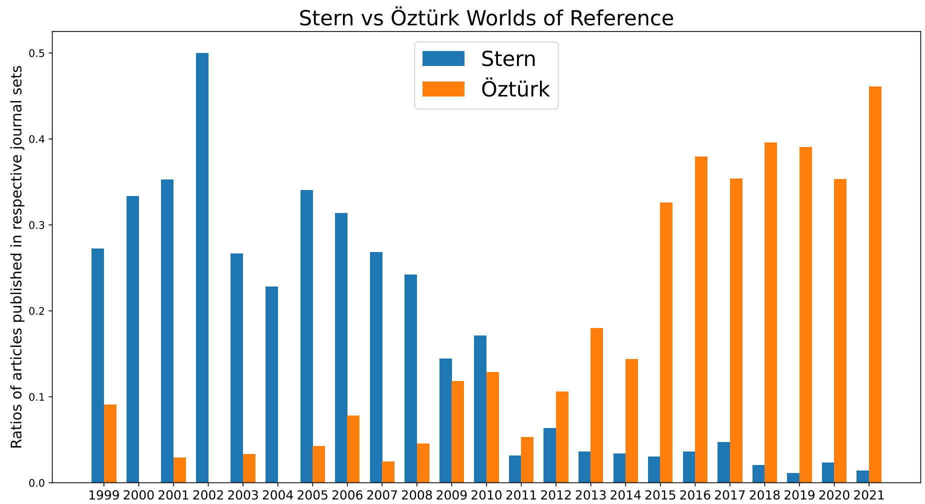

Figure 4: Relative number of papers within Stern's and Öztürk's worlds of references. The chart starts in 1999 as there are very few papers before that year and they are all published in Stern's world.

dicates 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. Papers in journals on top, mostly environmental economics journals, cite far more Stern than Öztürk, whereas papers on the bottom cite far more Öztürk. Remarkably, 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 as the set of journals in which at least one paper cited by him has been published, 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 4 is even more striking. For each world of references, we compute the yearly distribution of the ratio of papers published in each world of references relative to the total number of papers published that year. We can map the sliding of the field as more and more papers 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* 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. Table 9 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

Block	"Stern journals"	"Öztürk journals"
A	.429	.083
B	.159	.170
C	.000	.627
D	.008	.564
E	.022	.430

Table 9: Ratio of articles published in Öztürk and Stern "world of references" journals for each block.

do not publish in Stern journals and 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 epistemic cultures 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 – at the very least, we observe a trend towards specialization, while the relatively balanced share of positive vs. negative results indicates ongoing debate on the empirical materiality of the EKC. 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, papers 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. These cautious 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-report significant results vs. null results. 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 paper and what is reported in the abstract. The abstract may report positive results whereas the paper actually finds mixed results. A somehow different example of reporting bias is found in Apergis and Ozturk (2015) (cited 187 times, 9th most cited paper). The abstract reports "empirical support to the presence of an Environmental Kuznets Curve hypothesis". However, the paper investigates the EKC thanks to a regression involving up to the third power of GDP. 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 paper offers no explanations of why the third power has been overlooked.

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 in the field. He has been very influential from the inception of the field and it is therefore interesting to compare his views on the EKC for CO_2 to the kind of conclusions that the algorithm points to. Let us take its recent review as a starting point: Stern (2017) is quite critical of the EKC hypothesis. He shows raw data on emissions of CO_2 per capita and GDP per capita and concludes "there is little sign of an EKC effect" (p. 18). Following the argument already developed in Stern (2010), he estimates that "the relationship between the levels of [...] CO_2 emissions and income per capita is monotonic when the effect of the passage of time is controlled for". Any possible decreasing trend is thus attributed to a time-effect.

This skepticism towards EKC reflects a long-standing position. Indeed, the first paper (Stern et al., 1996, which was not concerned by CO_2) 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 papers 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 papers 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 papers that found an EKC for CO_2 rarely considers a bivariate relationship between CO_2 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 CO_2 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 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. All the perpetual innovations and additions in the field thus have stretched the object in different directions, to the point of strongly diminishing its consistency.

This is to say nothing of the multiple methods to assess its existence. For example, for the already discussed study on EKC for CO_2 in Vietnam, Al-Mulali et al. (2015) did not confirm the EKC hypothesis, because when the elasticity of CO_2 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. This is just one example of the traps that can be encountered in the field. And these seem to multiply, as, although the discussion has mainly been conducted with various econometric tools, more exotic methods are now also found, whose validity can be questioned.

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 papers, yet may diverge in what is meant by it. Human knowledge of a field and, more precisely, the justifications put forward by Stern in his assessment of the field, are important as they can inform us on what we cannot see from our perspective of external observers. In general, David I. Stern does not find the results obtained in the field convincing because "the naïve econometric approaches used in much of the literature are also problematic". It is certainly an impression that one can easily arrive at while reading some of the papers of our corpus, but it goes beyond as even a properly written paper may employ flawed econometric tools as was pointed by Wagner (2008, 2015). It is however very difficult to objectify without being an expert of each of the techniques employed in the literature and, in our perspective, it is outside of the scope of this paper to develop systematic rules to appraise the methodological aspects of papers on an automated basis. In a way, the evolution that we observed on figure 2 may also be framed as modest: in a nutshell, it remains 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 "EKC". We leave to further research the possibility of characterizing in more details 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. The use of semantic hypergraphs for the former 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. 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 SOx and NOx. The second wave and epistemic community, centered around 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. Notwithstanding, we observe on the whole that the share of positive results as reported in abstracts has consistently increased over the years, yet remains of the same order of magnitude as negative results — the debate is not closed. Furthermore, our appraisal of positive vs. negative results is based on authors' reporting in abstracts: casual examination of paper contents reveals that some positive abstracts correspond to more nuanced results in the paper 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.

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.

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.

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.

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.

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.

Dinda, Soumyananda. 2004. Environmental Kuznets Curve Hypothesis: A Survey. *Ecological Economics*, *49*(4): 431–455.

Grossman, Gene M. and Alan B. Krueger. 1991. *Environmental Impacts of a North American Free Trade Agreement*. Technical report, NBER, Cambridge (Ma.). Working Paper 3914.

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.

World Bank
. 1992. *World Development Report 1992. Development and the Environment*. Technical report, World Bank, New York.

Holland, Paul W, Kathryn Blackmond Laskey, and Samuel Leinhardt. 1983. Stochastic blockmodels: First steps. *Social networks*, *5*(2): 109–137.

Karrer, Brian and Mark EJ Newman. 2011. Stochastic blockmodels and community structure in networks. *Physical review E*, *83*(1): 016107.

Kuznets, Simon. 1955. Economic Growth and Income Inequality. *The American Economic Review*, *45*(1): 1–28.

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.

Panayotou, Theodore. 1993. *Empirical Tests and Policy Analysis of Environmental Degradation at Different Stages of Economic Development*. Technical report, International Labour Office, Geneva.

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.

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.

Shafik, Nemat and Sushenjit Bandyopadhyay. 1992. *Economic Growth and Environmental Quality: Time Series and Crosscountry Evidence. Background Paper for the World Development Report 1992*. Technical report, The World Bank, Washington.

Shahbaz, Muhammad and Avik Sinha. 2019. Environmental Kuznets curve for CO2 emissions: a literature survey. *Journal of Economic Studies*, *46*(1): 106–168.

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.

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.