Schmal, W. Benedikt

Vice versa: The decoupling of content and topic heterogeneity in collusion research

Journal of Economic Surveys

Provided in Cooperation with:
John Wiley & Sons

Suggested Citation: Schmal, W. Benedikt (2023) : Vice versa: The decoupling of content and topic heterogeneity in collusion research, Journal of Economic Surveys, ISSN 1467-6419, Wiley Periodicals, Inc., Hoboken, NJ, Vol. 38, Iss. 5, pp. 1686-1730, https://doi.org/10.1111/joes.12600

DOI: 10.1111/joes.12600

ARTICLE

JOURNAL OF ECONOMIC SURVEYS WILEY

Vice versa: The decoupling of content and topic heterogeneity in collusion research

W. Benedikt Schmal

Walter Eucken Institute, Freiburg (Br.) & DICE, Heinrich Heine University, Düsseldorf, Germany

Correspondence
W. Benedikt Schmal, Walter Eucken Institute, Goethestr. 10, 79100 Freiburg, Germany.
Email: schmal@eucken.de

Funding information
Deutsche Forschungsgemeinschaft, Grant/Award Number: 235577387/GRK 1974

Abstract

Collusive practices continue to be a significant threat to competition and consumer welfare. It should be of utmost importance for academic research to provide the theoretical and empirical foundations to antitrust authorities and enable them to develop proper tools to encounter new collusive practices. Utilizing topical natural language machine learning techniques allows me to analyze the evolution of economic research on collusion over the past two decades in a novel way. It enables me to review some 800 publications systematically. I extract the underlying topics from the papers and conduct a large set of uni- and multivariate time series and regression analyses on their individual prevalences. I detect a notable tendency towards monocultures in topics and an endogenous constriction of the topic variety. In contrast, the overall contents and issues addressed by these papers have grown remarkably. This caused a decoupling: Nowadays, more datasets and cartel cases are studied but with a smaller research scope.

KEYWORDS
antitrust, collusion, neural networks, systematic review, topic modeling

JEL CLASSIFICATION
D83, I23, L40, O33, Z13

1 | INTRODUCTION

In industrial economics and antitrust, cartels and collusive behavior seem to have gone off the radar somewhat. Even though the economic harm caused by it is non-negligible. Levenstein and Suslow (2006) report an 80% price increase caused by the international tea cartel and monopoly pricing of the US beer and the British ocean shipping cartel. Meta studies find an average price overcharge of 30%–50% and a median overcharge of 15%–25% by cartels in general (Bolotova, 2009; Boyer & Kotchoni, 2015; Connor, 2008). A proxy for the impact of collusion is the fines imposed: From 2000 to 2022, the European Union (EU) handed out penalties of some $31 billion.[1] Moreover, the EU plans to lose restrictions on anti-competitive behavior to promote the dissemination of environment-friendly technology (Schinkel & Treuren, 2021). Thus, collusion is and will continue to be a major issue for consumer welfare and antitrust policy.

This paper investigates research on collusion during the past two decades applying topical artificial intelligence tools. It reveals crucial changes over time and hints at developments in the literature that will likely affect competition authorities. Methodologically, I provide a novel way of understanding how collusion is addressed and identify the centers of academic attention. I compute latent topics by applying a structural topic model, a state-of-the-art natural language processing technique that relies on unsupervised machine learning. Hence, the paper brings together the economics of science, the history of economics, and antitrust research by providing a novel quantitative comprehension of the development of research on cartels for economists as well as antitrust scholars. Since industrial organization research is a significant source of inspiration for policymakers, legislators, and antitrust authorities, understanding the patterns of collusion research may directly impact the evolution of related public policy.[2]

The core question the paper aims to answer is: How did economic research on collusion among firms change in the past two decades contentwise, and how is this entangled with publication and citation patterns? By using machine learning techniques, I cannot only excerpt the whole strand of this literature. I also quantify developments over time that could be captured only descriptively otherwise. The main idea of natural language processing is the quantification of text data. Figure 1 illustrates this. This "word cloud" depicts frequent terms, where a larger font size implies a higher frequency. Here, I display words from topic 16 on leniency. It is important to note that I labeled the topic ex-post "leniency" as the unsupervised algorithm has "detected" it independently. Most words are shortened to their stems to account for inflections and variations between verbs and nouns.

Unsurprisingly, terms such as *leniency*, *effect*, *fine*, or *program* form the center of collusion research on leniency schemes. I quantify these (and all other terms) using a "structural topic model." While I explain the method in more detail later on, the core advantage is that it can incorporate explanatory variables to detect latent underlying topics in texts.I include time and place of publication as major determinants of which kinds of topics to expect. For every publication, I am able to compute the probability of each topic occurring. I utilize these topic probabilities to investigate topic prevalences over time using uni- and multivariate time series analyses. I amend my analysis with regressions investigating the relation between topics and citations.

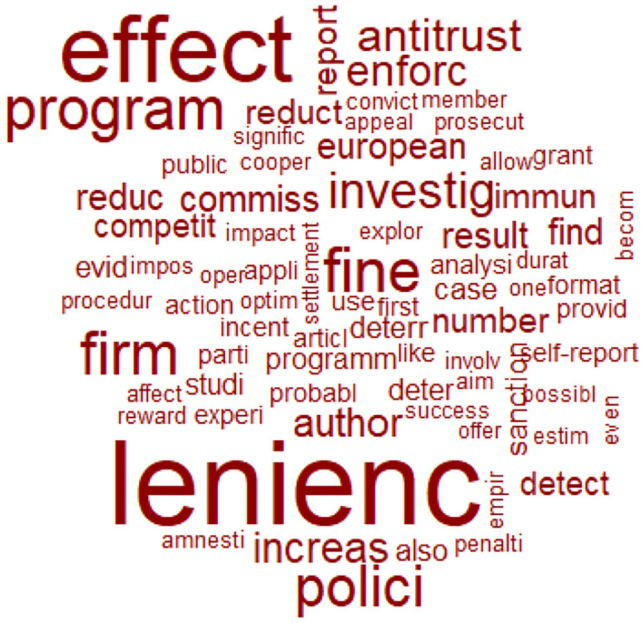

FIGURE 1 Wordcloud based on the computed latent topic on leniency. [Colour figure can be viewed at wileyonlinelibrary.com]

To conduct the aforementioned text analysis, I build upon 777 publications in leading economics and antitrust journals addressing collusion. Among this arguably full body of research in this sub-discipline, I identify 21 latent topics, analyzing the content of the papers with a structural topic model. I cluster these topics and assign them to broader categories. Based on that, the first part of the subsequent analysis studies the evolution of topic categories over time. To achieve this, I extract the predicted probability for every topic to occur in a specific paper. It becomes visible that competition analyses studying the generalized behavior of firms are in retreat. Especially in recent years, empirical case studies of disclosed cartels and their court cases have seen a considerable boost. Furthermore, empirical evidence exists for specialization, that is, the focus of a paper seems to be narrowed to fewer topics.

The case study-related topics are strongly correlated with the leading so-called "top 5" journals in economics. While other general interest and broader field outlets cover various topics, specific IO journals also seem to narrow their focus to a smaller set of topics. This may end up in an endogenous process of streamlining research. Conducting a repeated cross-section regression analysis on the relationship between citations and topics, I do not find any significant positive relationship but three topics to be negatively correlated with citations per year. Remarkably, these topics are in decline over time in terms of the predicted probability to occur in a paper. A subject-wise decrease does not match this content-wise decline: Applying the machine learning technique of neural networks, I can show that the superficial *subjects* of papers become more diverse over time while the underlying actual considerations are more closely tied together.

In the last step, I conduct a multivariate time-series analysis to understand how the decline in competition analysis and the upswing in case studies are intertemporally related. The latter do not supersede the former but fill the gap the reduction in the latter leaves. The discipline tends to have lost some interest in the stylized models of markets and instead turned towards case studies. It serves as additional evidence for endogenous constriction of academic research on collusion, which may affect the ability of competition authorities to chase and disclose collusive structures.

This paper is a systematic literature review based on natural language processing of industrial organization research. It contributes to the economic strand of the science of science literature that aims to understand how the process of creating and disseminating scientific progress works. Furthermore, it also reflects the rules and habits of the discipline of economic research. Earlier research by Einav and Levin (2010) discussed the developments in empirical IO, while Hovenkamp (2018, 2021) recently criticized the current state of antitrust research. Lancieri et al. (2022) find evidence for a decline in antitrust enforcement in the United States but hardly any evidence for antitrust research playing a role in that. Shapiro (2020) raises the concern that antitrust economics has focused on technical aspects of markets but rather misses out on the broader economic issues affected by, for example, higher prices or less innovation. Jones and Kovacic (2020) discuss the problem of antitrust authorities dealing with the market power of digital platforms.

This paper aims at contributing a detailed understanding of research on cartels. Methodologically, it follows Qasim (2017), who analyzes interdisciplinary sustainability research, and Ambrosino et al. (2018), who already clustered economic research in total via topic modeling. Natural language processing recently gained popularity in economic research (Gentzkow et al., 2019). Larsen (2021) uses it to understand market uncertainty. Much seminal economics research applying topic modeling investigates central bank communication and its macroeconomic effects (Hansen & McMahon, 2016; Hansen et al., 2018, 2019). They use the latent Dirichlet allocation (LDA) algorithm. While easy to use, structural topic modeling (STM) allows for a more granular analysis with less restrictive model assumptions. Therefore, I choose the latter technique. The STM model begins to gain traction in economics and business research (see, e.g., Ebadi et al., 2021; Ferrara et al., 2022; Kumar & Srivastava, 2022; Moschella et al., 2020; Nathan & Rosso, 2022; Schmal, 2022b)

Last, Ash et al. (2022) have shown that economic ideas can have a high impact on practitioners as they show that economics classes for judges influenced their verdicts. An important recent example from antitrust is the OECD (2017) policy paper on algorithms and collusion, which draws from recently published economics and antitrust papers such as Ezrachi and Stucke (2016), Mehra (2015), Green et al. (2014), Potters and Suetens (2013)—even though such a topic almost naturally heavily relies on websites and online media.

The remainder of the paper is structured as follows. Section 2 describes the data and the machine-learning-based structural topic modeling methodology. Later on, it also describes the econometrics behind the analysis and how I cluster the computed topics. Section 3 presents the paper's findings, in particular, the time series and regression analyses. Section 4 contextualizes and concludes.

2 | EMPIRICAL APPROACH

2.1 | Data

The foundation of the analysis is the data from the Scopus database, as shown in Figure 2, obtained with the Python API *pybliometrics* (version 3.3.0, Rose & Kitchin, 2019). I gather 34,564 publications in 35 journals from 2000 to 2021, downloaded from April 4–6, 2022. I identify papers that address collusion by choosing those with at least one of the operators "collusion," "collusive," "cartel," "bidding ring," or its declensions in its title, abstract, or keywords.[3] By that, the number of relevant papers collapses to 777.[4]

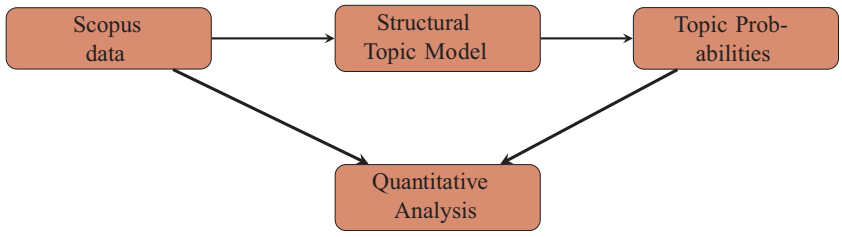

FIGURE 2 Data used in the quantitative analysis. [Colour figure can be viewed at wileyonlinelibrary.com]

I exploit the abstracts of papers published in leading IO and antitrust outlets and their related metadata. While a full-text analysis would encompass more information, abstracts are focused on the core techniques and messages a paper applies and wants to communicate to the reader. Due to its brevity, an abstract is usually not decorated with illustrative examples not used for deeper analysis, as often seen in introduction sections of full papers. This helps separate important from less important cases in the text analysis and follows the approach of Wu et al. (2019) and Gropp et al. (2016). Sybrandt et al. (2018) find that full texts may provide more information but, in turn, more "intruder" terms that obfuscate the topic modeling and subsequent interpretation.[5]

I preprocess the text data from the abstracts to ease the analysis. It includes the removal of punctuation, transforming all words into lowercase terms, and removing stopwords, for example, conjunctions or articles that have no real meaning but are necessary for humans to read and understand a text.[6] I manually remove so-called "boilerplate" words (see Lämmel & Jones, 2003) that are technically part of the abstracts but do not contain any useful information, such as copyrights of the publishers. Furthermore, I remove the operators "cartel" and "collusion" as these terms, by definition, should be part of virtually every abstract. Hence, they add no information on underlying topics *within* the subset of papers about cartels and collusion.[7]

Last, I stem the words, that is, plural forms or the past tense for verbs disappear. For example, "cooperation," "cooperate," and "cooperative" collapse to "coopera." This bears the loss of some information, but it comes with a considerable gain in analytic power, as otherwise, every version of a word would be considered an independent term.

2.2 | Structural topic modeling

Based on the body of papers, I conduct natural language processing. I work with the papers' abstracts to understand deeper relations and developments in collusion research over time. To do that, I apply the technique of topic modeling. This method was made popular by Blei et al. (2003). The core idea is as follows: A number of D text elements (usually called documents, in my case, paper abstracts) contain K "topics," that is, clusters of words that belong together.[8] Each document d consists of a number of W words.[9]

The ultimate goal is identifying a pre-defined number of K latent topics without a prior which words belong to them. A data-generating process for each document is assumed in which every document consists of several topics, which themselves consist of a set of words that have their own probability of belonging to a particular topic. In the following, I sketch the mechanism of this unsupervised machine-learning algorithm based on the work of Blei et al. (2003) and Roberts et al. (2013). Note that I mostly follow the notation of the latter even though I adjust parts to address an economics readership better.

I apply the *structural* topic model (STM, Roberts et al., 2013) using the R package *stm* of Roberts et al. (2019) because it allows for the incorporation of covariates as independent variables for the prevalence of topics. Based on that, I estimate regressions on relationships between specific topics and covariates X_d. Every document has a set of topics in it. The STM assumes that there is an individual prior for the particular set of topics for each document, which depends on a vector of covariates. Put differently, the covariates are assumed to determine which topics appear in a particular paper. Thinking of academic publications, a paper's outlet is likely to be related to the topics as the *Journal of Economic Theory* may cover different topics than *Empirical Economics*.

Second, a vector $\theta_d \sim logit\mathcal{N}(X_d\gamma_k, \sigma)$ for each document contains the document-specific probability of each topic to occur there. γ_k is a vector of covariates for X_d, with $\gamma \sim \mathcal{N}(0, \sigma_k^2)$, such that one assumes as a prior the document-covariates to be uncorrelated with the topics and only deviate from this assumption for a strong correlation between the two. Each document has a document-specific probability for a word to occur. It depends on the overall distribution of words (m) and the covariates (X_d). It is captured by $\beta(X_d, m)$ and describes the probability of a particular word occurring in a particular document, given that it contains some particular topic. β is a matrix of size $K \times W$, that is, it contains for every word ($w_1 \dots w_N$) and every possible topic ($k_1 \dots k_N$) the probability to occur.

Based on θ_d, one draws the topic ($\hat{k}_w$) that is most likely for every word w in every document using a multinomial logit model, that is, $\hat{k}_w \sim Multi(\theta_d)$. Last, one draws an actual word $\hat{w}$ that has the highest probability to occur given $\hat{k}_w$ and β, that is, $\hat{w} = max\{p(w_n|\hat{k}_w, \beta)\}$ while the actual words are meant to be distributed $w_{d,n} \sim Multi(\beta_d^{k=\hat{k}_w})$). One can easily see that $\hat{w}$ should converge towards w. As the actual topics within each document are latent, structural topic modeling uses an expectation-maximization algorithm that iteratively tries to find a local maximum likelihood given the prevalence of latent topics.

Determining the number of underlying topics K in a body of documents is a crucial task. There exists no "one size fits all" approach to elicit the optimal K^* (see, e.g., Grimmer & Stewart, 2013; Wehrheim, 2019). Nevertheless, several key measures contain important information on the quality of the topics dependent on the number of topics. One key measure is semantic coherence, which states that words that occur very frequently within topic k_i should also occur in a document, given this document contains k_i (Mimno et al., 2011). While this is a convincing concept, it suffers from the problem that for low K, semantic coherence is high by construction.

As an antagonistic metric, I use exclusivity as proposed by Bischof and Airoldi (2012). The core idea behind that is that words that are only in one topic very frequently are "exclusive" while words with relatively equal frequencies across topics are somewhat non-exclusive. Consider the word "collusion" in the given context that should occur in many topics. As this is such a big keyword—it is part of nearly all abstracts by construction—I exclude it from the documents. Words that occur in many topics should lead to a higher semantic coherence because, in this case, the topics are more likely to cover many words in the actual abstracts. In contrast, these topics should have a low exclusivity, so the trade-off between these two measures should provide a sufficient choice of K. In Appendix B, I sketch in more detail the measures I compute to choose K. Overall, here I conclude to set $K^* = 21$.

$$\theta_d \sim year_d + journal_d$$

I use a structural topic model, so I am able to add covariates to the topic regressions, which, in general, are fitted as linear models with the expected topic prevalence as the dependent variable.

Accordingly, θ_d represents the expected prevalence of the 21 topics. I include *year* as a temporal variable to account for changes over time (linearly without any spline). I add the *journal* as a significant covariate that should be related to a paper's particular content. To draw more general conclusions and increase statistical power, I group the journals into five categories, as shown in Table A.1 in the appendix, because clusters of journals with similarities in scope and methods exist.[10] I later use these broader categories for the correlations between journal clusters and topics. Further variables could be added, but due to the model structure, only highly relevant covariates get assigned a value different from zero, which suggests a small set of important variables.

Most parts of this paper rely on a structural topic model. I amend my analysis with a machine learning algorithm that relies on neural networks. Other than the statistical expectation-maximization approach taken in the STM, neural networks set up a plethora of interrelated nodes, to some extent equivalent to neurons in the human brain. By that, a neural network-based model draws relations between variables or topics from the relations between nodes and the particular weighting of these relations (see, e.g., Amari, 1995). In the present case of text analysis, I use the *doc2vec* neural network algorithm developed by Le and Mikolov (2014). The unsupervised algorithm autonomously learns relations between words in a text (here in academic publications) and converts it into a lower-dimensional vector for each document. Based on that, I compute the inner product of these vectors for each pair of publications within a year to measure the similarity of these papers. I average all similarities per paper and in a second step per year to compare changes over time.

2.3 | Quantitative estimation

The predicted topic prevalences drawn from the structural topic model and its unsupervised machine learning approach build the foundation of the analysis. As sketched in Figure 2, these probabilities are amended with additional publication data from the Scopus database. To analyze the dynamic dimension of topic prevalences, I construct multiple time series based on annual topic averages and conduct various stationarity tests, particularly the widely applied augmented Dickey-Fuller (ADF) test and the Phillips-Perron (PP) test. Both procedures share the weakness of non-stationarity being the null hypothesis. That is, no rejection of the null does not imply the existence of a unit root. Kwiatkowski et al. (1992) suggest the KPSS test that puts non-stationarity as the alternative hypothesis, such that a rejection of the null should be more unambiguous evidence for a unit root. I use all three test procedures to ensure valid findings.

ADF and PP tests come in three different types, namely a linear model with no trend and no drift (type 1), one with no linear time trend but a drift (non-stationarity in the mean; type 2), and one with a linear trend and a drift (non-stationarity in mean and variance; type 3). Additionally, the number of lags in the auto-regressive process can be specified. While annual data usually apply one lag as it mostly relies on the previous year, I choose a lag of $n = 2$. It has both an economic and a technical foundation that coincide: Economic publishing faces a notable lag between the submission of a paper and its publication.[11] A two-year time lag is more reasonable, in which the topic prevalence of year t is a function of $t - 2$.[12] The range of my data implies $max(t) = 22$, which leads just as the theoretical considerations to $n^* = 2$. For the unit root tests, I also rely on the test statistics of the second lag.

Next, I conduct a repeated cross-sectional analysis of the relationship between latent topics and the reception of papers with these topics in the literature. My dataset allows me to look at the number of citations per paper (until 2022). I am able to study the relationship with the presence of specific topics. I use the number of citations per year of a specific paper i as the dependent

variable and the expected topic probability T as an explanatory variable. This looks as follows:

$$f(c/y)_i = \beta_T T_{ij} + \beta_{OA} \mathbb{1}_{OA} + \beta_A A_i + Y \times J + \epsilon_i \,, \tag{1}$$

whereby β captures the estimated coefficients of the covariates. Citations per year on the LHS are transformed to their logarithmic form. Figure A.5 in the appendix shows that this measure is approximately normally distributed. However, since there are papers without citations, some 13% of the observations get lost. Hence, I apply two transformations to avoid reducing the already rather small sample size. First, I use $log(c/y + 1)$ to avoid the dropout of zero values. Second, I apply the widely used hyperbolic sine transformation, that is, $log(c/y + \sqrt{c/y^2 + 1})$. Solely relying on the "pure" logarithm would not be sufficient as papers with zero citations are an important part of the analysis, as not being cited is relevant information, too. On the right-hand side, I add the logarithm of topic j to understand how the prevalence of this particular topic corresponds to citations.

I include two crucial control variables: First, the corresponding author (A_i) to account for the fact that a paper's visibility and reputation certainly depend on its author(s).[13] Second, I use a binary dummy for open access ($\mathbb{1}_{OA}$) to capture the fact that open access tends to affect citations, even though the evidence is mixed (McCabe & Snyder, 2021, 2014; Gaule & Maystre, 2011). I also add a *year×journal* fixed effect to account for changes over time and journal reputation, as it is crucial for visibility, credibility, and citations.

Last, I exploit the fact that we have several topics co-moving, such that one can set up a multivariate time series analysis, that is, a vector autoregression (VAR) model. I estimate a model including several groups of topics, which I name *categories*,[14] and usually apply a lag length of two periods. It is reasoned by the publication lag in economics as well as econometric tests for the optimal VAR length. Details are specified and explained in the results section for each estimation.

Based on all these computations, I obtain essential insights into the dispute and contention about collusive behavior within the scientific community, the evolution of approaches over time, and the appeal of particular topics captured by relationships between topics and journal types as well as by the "citability" of particular topics.

2.4 | Clustering topics

From a statistical and economic perspective, it is more useful to cluster and aggregate topics instead of studying them individually. However, doing so naturally raises the question of how to do that.[15] I choose to adopt the Institutional Analysis and Development (IAD) framework, originally developed by Elinor Ostrom (see for an overview, e.g., McGinnis, 2011; Ostrom, 2011) to structure and understand jointly used or managed common pool resources, it has found plentiful applications in political science but less so in economics. However, Schmal (2022a) indicates how far cartels can be understood to generate "common pool resources" on which the Ostromian IAD framework can be applied to understand their functioning better. Based on that, I argue that cartel research can be clustered in that way as the topics should reflect the actual structure of cartels. Nevertheless, alternative approaches and methods of clustering would be conceivable as well, just as mentioned beforehand.

The framework is shown in Figure 3. It consists of three major blocks: On the left, there are topics that address issues outside of cartels themselves. Here, I go—equivalent to the original model—from broad to narrow, starting with topics on the overall economic environment to

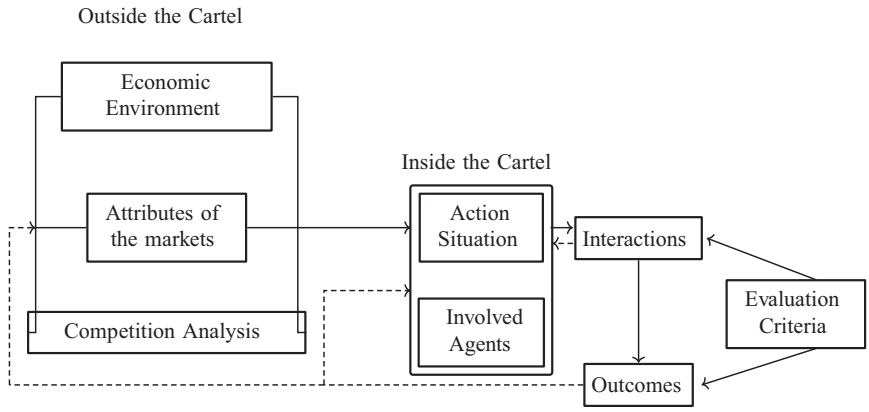

FIGURE 3 The structure of the IAD framework adapted for clustering. Original taken from (Ostrom, 2005, p. 15), see Figure A.1 in the appendix.

attributes of specific markets. The bottom item is the competition analysis that often studies partial equilibrium effects of slight variation in parameters relevant to firm competition, in my case, also linked to collusion.

In the middle, topics on issues within cartels and collusive agreements find their place. This could be, for example, the effect of leniency schemes that allow involved firms or actual agents, such as managers, to leave a cartel without punishment in case they are the first to report the cartel to authorities and to provide them with substantial information about it. On the right, there are "evaluation criteria" devoted to assessing the interactions between the involved agents and the outcomes of cartels, such as firm profits or changed market characteristics. This links back to the pillar on issues outside a cartel on the left.

The block "inside the cartel" consists of the involved agents as well as the situation in which they find themselves. I consider this to be the part in which no stylized agents or firms act but employees who consider their individual constraints besides incentive constraints for profit maximization of their respective firms. It is a less addressed issue in industrial organization research and not at the center of the following analysis. I attempt to map the underlying topics of existing antitrust and collusion research to this framework to get a first and more detailed understanding of where research on antitrust and collusion stands and what strands might be less investigated up to now.

Such a pre-defined model provides a fixed frame for clustering and avoids an ex-post fitting. Nevertheless, assigning the computed topics to a specific component might object to subjective sorting. With a rigorous explanation and an extensive definition of the components, I aim at avoiding arbitrariness. As such, I follow the suggestion of Cherrier and Svorenčík (2018) to combine quantitative methods with qualitative techniques—as done with this clustering—when studying the history of economics as a discipline.

3 | RESULTS

3.1 | Latent topics and topic mapping

In the first step, I compute the latent topics and their respective content based on the optimal K^*. To elicit the terms related to each topic, I use two metrics. First, I show the words with the

highest probability of appearing, given that a document contains a certain topic. Second, I use the FREX measure as proposed by Airoldi and Bischof (2016) that combines the frequency (FR) and the exclusivity (EX) of a specific word in one measure, such that I provide those words with the highest "FREX" value for each topic. It is important to remember that the keywords are word stems. It circumvents the issue of declensions and focuses on the core content but misses linguistic details. I obtain a set of $K = 21$ latent topics as shown in Table 1—together with the most important keywords of each topic.[16]

Table B.2 in Appendix B shows the correlations between topics and highlights that there generally exists a low level of correlation, which is mostly slightly negative. Thus, there is no data-driven pattern on which topics to aggregate in clusters. It re-emphasizes the importance of manual clustering based on an informed assignment. I manually assign a label to each topic to make it easier to follow the content of the respective topics. The numbers do not imply any ranking but should be considered an index. Next, I map these topics to the framework as sketched in Figure 3 before. Table 2 shows this.

While topics 1 and 2 address case studies and court cases, they cannot be located directly within the framework. The best fit may be the "evaluative criteria." Such papers are often extensive quantitative analyses that implement structural estimation methods. Topic 3, however, discusses market shares arising from competition, cooperation, or alliances. This can be related to the "outcomes" of the decision situation, that is the pillar on the right side graphically in Figure 3. The main topic keywords hint at industry and competition characteristics in a way that fits actual general market outcomes. Topic 4 on auction theory can be linked to the competition analysis within the "outside the cartel" section, as the design of procurement auctions shapes the participants' behavior. Topic 5 on information mechanisms fits into the action situation and the related information of the participants and how they are exploited.

Topic 6, again, is a case study topic that captures actual cases. In IO research, such examples are natural starting points, as large blown-up cartels or mergers between leading firms often attract great attention among the informed public. Again, the "evaluative criteria" might capture this topic the most. However, it is a "meta" topic, such as topics 1 and 2. In contrast, topic 7 on principal-agent issues fits into the heart of every cartel, namely the action situation, and within that, it fits the participants, their positions, and their actions. Topic 8 deals with mark-ups and is, by that, an outcome of collusive activities. Topic 9 on law enforcement fits the interactions outside the action situation as it addresses the outcomes (overall profits depend on potential fines) and corresponds to exogenous variables such as the legal framework and the general attitude towards cartels in society and among politicians. The topic on capacity constraints (#11) seems to be very theory-based and abstract and, by that, tends to correspond to the field of competition analysis outside an actual cartel.

The topics on gasoline and airlines (#10 and #14) again tend to be case-study-based. This, however, is the case as these two markets have primary importance and are distinct as they sell relatively homogeneous goods (fuel and transport between two airports), which makes it particularly interesting to study. Studies in this domain usually address the market as a whole. Therefore, I assign it to the "attributes of a market" within the "outside the cartel" section.

The topic of game theory (#12) is more difficult to allocate. On the one hand, it addresses group behavior and how agents or players interact. However, the equilibrium focus of this topic (and game theory per se) seems to fit into the competition analysis, where it is heavily applied. One has to keep in mind that these formalized considerations suggest general behavior in the form of strategies to be optimally played. Mainly related to sophisticated solution concepts, it is

TABLE 1 Latent topics in joint 21st century antitrust and collusion text corpus.

Topic 1: Antitrust Overview	
Highest Prob:	competit, articl, author, enforc, antitrust, polici, develop
FREX:	author, remedi, jurisdict, sanction, articl, review, discuss
Topic 2: Court Cases	
Highest Prob:	court, law, case, competit, decis, articl, damag
FREX:	court, infring, suprem, litig, settlement, judgment, rpm
Topic 3: Market Shares	
Highest Prob:	market, firm, competit, cooper, model, share, price
FREX:	cooper, allianc, market, ventur, format, share, joint
Topic 4: Auction Theory	
Highest Prob:	auction, bidder, bid, first-pric, mechan, use, two
FREX:	bidder, auction, first-pric, ring, second-pric, bid, sealed-bid
Topic 5: Information Mechanisms	
Highest Prob:	inform, privat, signal, communic, equilibrium, mechan, game
FREX:	signal, inform, privat, payoff, communic, correl, type
Topic 6: Collusion Cases	
Highest Prob:	antitrust, econom, competit, athlet, case, analysi, law
FREX:	athlet, ncaa, colleg, ncaa, oil, restraint, wto
Topic 7: Principal-Agent Issues	
Highest Prob:	agent, contract, princip, incent, optim, deleg, organ
FREX:	princip, agent, deleg, corrupt, worker, supervisor, contract
Topic 8: Mark-Ups	
Highest Prob:	price, damag, use, overcharg, predict, guarante, subject
FREX:	guarante, overcharg, subject, damag, price-match, lysin, method
Topic 9: Law Enforcemenmt	
Highest Prob:	law, enforc, competit, state, signific, deterr, fine
FREX:	complianc, deterr, resourc, freedom, centuri, state, foreign
Topic 10: Gasoline Market	
Highest Prob:	price, retail, market, use, data, cost, model
FREX:	gasolin, retail, ceil, station, adjust, allianc, partner
Topic 11: Capacity Constraints	
Highest Prob:	firm, market, product, price, sustain, cost, profit
FREX:	firm, capac, integr, bertrand, sustain, invest, rival
Topic 12: Game Theory	
Highest Prob:	equilibrium, group, incent, game, player, threat, econom
FREX:	group, threat, borrow, equilibrium, ineffici, lend, game
Topic 13: Vertical Relations	
Highest Prob:	competit, supplier, manufactur, power, retail, suppli, effect
FREX:	manufactur, supplier, exclus, suppli, chain, ncaa, balanc
Topic 14: Airline Markets	
Highest Prob:	model, airlin, exchang, competit, bargain, contact, account
FREX:	airport, contact, airlin, exchang, bargain, japan, account

(Continues)

TABLE 1 (Continued)

Topic 15: Merger Analysis	
Highest Prob:	merger, coordin, industri, firm, use, effect, price
FREX:	merger, coordin, merg, margin, multiproduct, screen, activ
Topic 16: Leniency	
Highest Prob:	lenienc, effect, fine, polici, program, firm, investig
FREX:	lenienc, program, immun, fine, reduct, commiss, programm
Topic 17: Antitrust Enforcement	
Highest Prob:	competit, antitrust, use, market, manag, determin, effect
FREX:	ship, pool, manag, connect, patent, determin, bank
Topic 18: Welfare Analysis	
Highest Prob:	cost, welfar, trade, price, penalti, higher, consum
FREX:	penalti, welfar, link, qualiti, larger, differenti, good
Topic 19: Procurement Auctions	
Highest Prob:	bid, procur, use, test, construct, auction, model
FREX:	procur, construct, bid-rig, bid, statist, round, rig
Topic 20: Cartel Organization	
Highest Prob:	price, increas, cost, firm, market, profit, consum
FREX:	price, announc, discount, transpar, leadership, cycl, war
Topic 21: Market Entry	
Highest Prob:	entri, competit, effect, market, model, chang, arrang
FREX:	entri, arrang, entrant, varieti, barrier, chang, characterist

$K = 21$. Text corpus contains 777 documents (i.e., abstracts), 2,601 terms and 37,785 tokens. EM-Algorithm iterations: 75. Spectral initialization. The topic labels are assigned manually.

questionable whether managers involved in a cartel solve decision situations in such a way when interacting with their colleagues.

A similar issue is addressed in the following topic on vertical relations, particularly related to grocery retail stores or chains. Again, it addresses either competition analysis or "attributes of the market." In the case of firms, this might rather be the characteristics of the industry. Merger analysis (#15) is another core issue of industrial organization and antitrust in general. It may play a minor role in collusion research but is still non-negligible. It is also a topic that discusses market's economic environment and characteristics, as it usually studies how the overall market is affected following a successful merger.

Topic 16 on leniency addresses potential outcomes within the action situation, as leniency is the central "joker" for colluding firms and their involved agents to leave a cartel without being penalized.[17] This changes incentives and, hence, net costs and benefits for the participants. Content-wise, closely related to leniency is the topic of antitrust enforcement. Major terms are – similar to leniency – "fine," "criminal," and "manage," that is, terms that affect the action situation and potential outcomes, but even more the interactions without the actual action situation. This is why I assign it to the "interactions." Topic 18 is labeled "welfare analysis." If one looks at the most important words, it seems to capture welfare considerations within competition analyses. This is why I categorize it like that, even though the separate sub-category "economic environment" would also be conceivable. The topic on procurement auctions (#19) can—as the related topic on auction theory—be assigned to competition analysis.

TABLE 2 Clustering of the topics.

Outside the cartel	
Attributes of the market	Gasoline (10), Vertical Relations (13), Airlines (14), Merger Analysis (15), Market Entry (21)
Competition analysis	Auction Theory (4), principal-agent issues (7), Capacity Constraints (11), Game Theory (12), Welfare Analysis (18), Procurement Auctions (19)
Inside the cartel	Information Mechanisms (5), Leniency (16), Cartel Organization (20)
Interactions	Law Enforcement (9), Antitrust Enforcement (17)
Outcomes	Market Shares (3), Mark-Ups (8)
Outside the framework and case studies	Antitrust Overview (1), Court Cases (2), Collusion Cases (6)

Topic 20 on (internal) cartel organization seems to be one of only a few topics that directly fit into the action arena. Here, terms such as "price," "leadership," "transparency," or "announce" directly correspond to agents deciding and coordinating on how to price their products and whether to stick to cartel agreements in a certain period or not. The very last topic, #21, captures market entry. Again, this is an issue outside of an actual cartel. It might be endogenously driven by the cartel's decisions, for example, by raising prices so much that an external firm might be inclined towards entering this particular market. However, single participants are likely to be unable to drive the whole market in this direction. Hence, "market entry" captures an attribute of a market as entry barriers form a vital characteristic of a market.

As mentioned, three topics do not fit into the applied cluster, namely the overview topic #1 and those on particular collusion cases (#2 and #6). While the former topic is self-explaining, the case-based topics constitute an ex-post analysis of existing cartels. One might also question the distinction between the competition analysis and the so-called "action arena" as topics like leniency can be analyzed theoretically and enter the decision-making process ex-ante. The assignment, however, considers that the three topics in the action arena category apply to decision-makers within the actual situation. While incentive constraints for a firm always apply and can be computed in advance, the decision to become a whistle-blower or to hold back information is made by being involved in the collusive agreement as a manager. Therefore, it directly relates to the actions of the acting individuals.

3.2 | Topic prevalence over time

Looking at the topic clusters shown in Table 2, most topics address issues "outside the cartel." However, it does not tell anything about the share of this cluster in collusion research. This section changes it by investigating the expected probabilities of topics aggregated to the clusters shown beforehand and per year to look at developments over time. Figure 4 shows the aggregated expected topic probabilities as assigned to the components above. It is done by predicting an expected value for each of the 21 topics per paper and taking the means per year.[18]

For many years, topics "outside a cartel," such as the competition analysis, counted for more than 50% of the expected topics. Even though there is a slight downward trend, one cannot say that

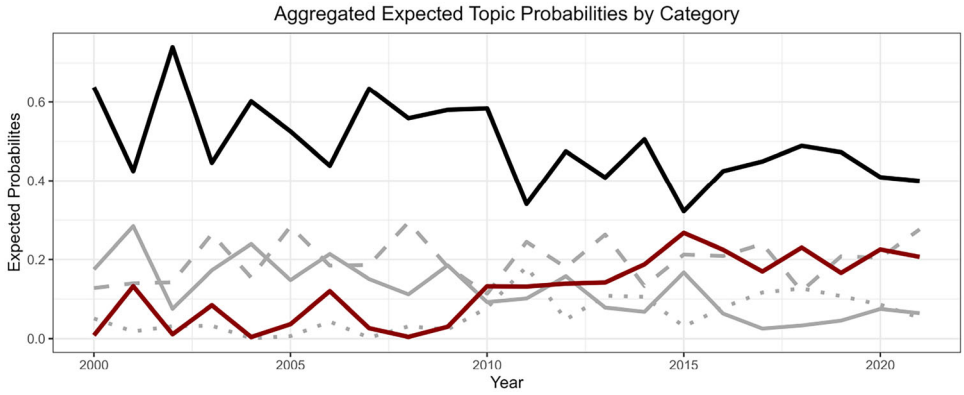

FIGURE 4 Topic probabilities over time. Black solid: Outside the Cartel (Economic Environment, Attributes of the Market, Competition Analysis); gray dashed: Inside the Cartel; gray dotted: Interactions; gray solid: Outcomes. Red solid: Case studies outside the framework.
[Colour figure can be viewed at wileyonlinelibrary.com]

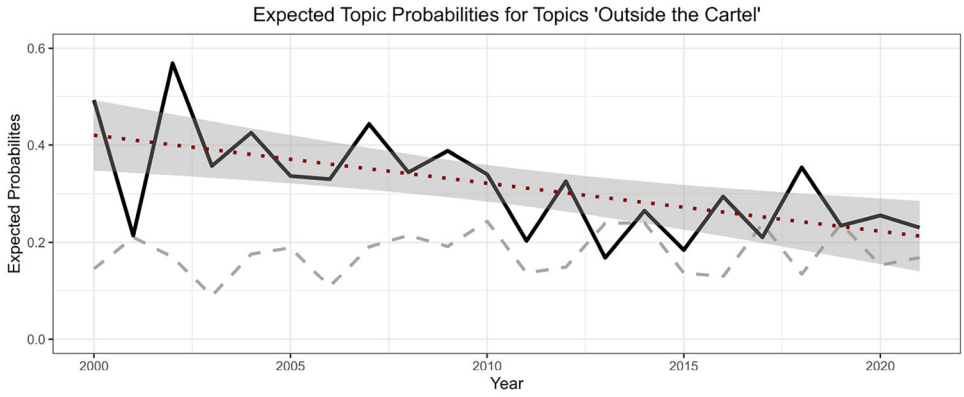

FIGURE 5 Topic probabilities of the exogenous variables. Black solid line: Competition analysis; Gray Dashed line: Attributes of the market. Red dotted line: Fitted linear regression line for the 'competition analysis' category.
[Colour figure can be viewed at wileyonlinelibrary.com]

another component would take the lead. In contrast, the prevalence of studies on within-cartel issues, the grey dashed line, seems to fluctuate around 20%. Interactions and outcomes appear relatively stable over time as well.

An important observation in Figure 4 is a drop in the share of the exogenous topics from 2010 to 2011. Figure 5 shows this in more detail by decomposing the "outside the cartel" block into its two components. One immediately sees that the drop in the overall category is mostly driven by the drop in "competition analysis" but accompanied by a drop in the expected probabilities of the "attributes of the market" component (gray dashed line) between 2010 and 2011, which does not seem to be part of a major drop but rather a temporal fluctuation. Therefore, the overall decline of the block appears exaggerated. Fitting a linear regression for $E[Pr(Comp.\ Analysis)]$ might be more informative and highlights the category's continuous decline over time.[19]

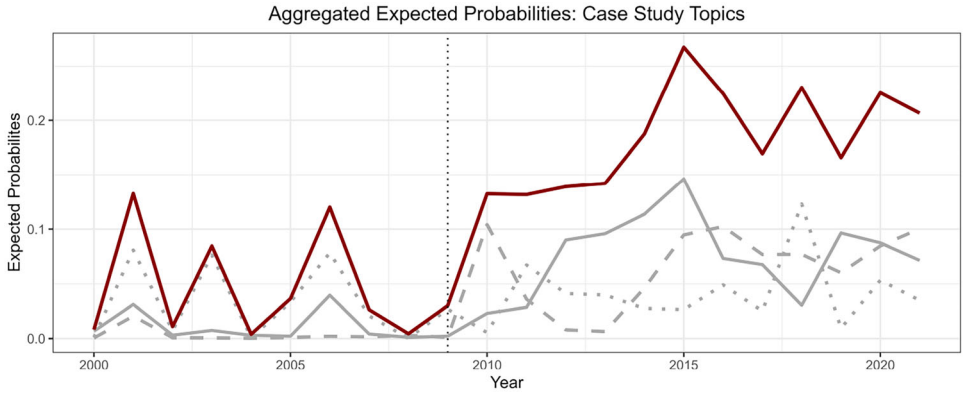

FIGURE 6 Topic probabilities of topics outside of the IAD framework. Red solid line: Sum of the three topics; Gray solid line: Topic 1—Antitrust Overview; Gray Dashed line: Topic 2—Court Cases; Gray dotted line: Topic 6—Collusion Cases.
[Colour figure can be viewed at wileyonlinelibrary.com]

To ensure that this decline is not misinterpreted, or put differently, to ensure there exists significant non-stationarity in the aggregated topics of the "competition analysis" subcluster within the "outside the cartel" block, I conduct stationarity tests as described in Subsection 2.3. The detailed results for all three tests can be found in Table A.4 in the appendix. The ADF test cannot reject the null hypothesis of non-stationarity for one or two lags for types 2 and 3. The PP test, however, rejects the null of non-stationarity for these two types. As Leybourne and Newbold (1999) have shown, the ADF and PP tests are susceptible to diverging, especially for AR(2) specifications. The KPSS test supports the economic hypothesis of non-stationarity by clearly rejecting the null of stationarity for a lag parameter of 2 ($p = .0182$).

In contrast to the "competition analysis" topics, one can observe a remarkable surge in the prevalence of case-based topics outside the framework. It is shown in Figure 6. While there seems to be in the aggregate of these topics some major fluctuations in the early 2000s, one can observe a steady rise from 2009 on, which peaks in 2015 and stabilizes at a high level around 20% afterwards. Furthermore, up to 2009, the prevalence of the aggregated three topics was mainly driven by collusion cases (#6), that is, a topic that mainly addresses leaked-out or blown-up cartels by authorities and their economic implications. From 2008 on, both rather antitrust-focused topics began to rise in their respective aggregated prevalence. As Table A.5 in the appendix shows, no ADF test rejects the null for two lags but the PP test does, again. The KPSS test rejects the null of stationarity, such that a unit root seems to be present. Complimentary, the KPSS test rejects the null of stationarity. Chow tests on the 5% significance level for a structural break in the data suggest that there might be one in 2009 ($F = 4.9398$, $p = .0195$). Overall, I am sure this surge is statistically supported.

It is likely to be driven by an intensive margin shift towards these topics and an extensive margin effect. Especially among antitrust journals, there are few publications addressing collusion in the early 2000s (see Table A.2 in the appendix). Second, both the number of cases and the number of fines awarded, for example, by the European Commission, only began to rise from 2005 and subsequent years on.[20] Additionally, the EU put a substantially revised leniency regulation in place in 2006 (see, e.g., Wils, 2007). It might have triggered new academic inquiries after disclosed cartels employed this leniency scheme. Nevertheless, the extensive effect also affects the

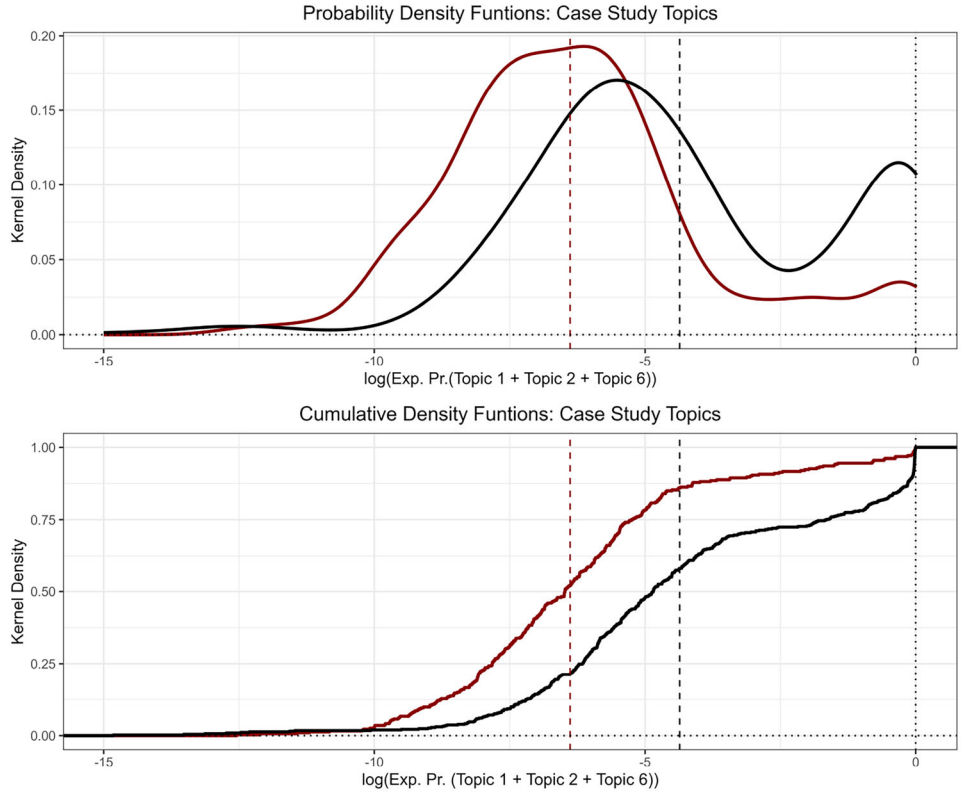

FIGURE 7 Topic densities of external topics prior- and post-2010. Density functions for the logged sum of external topics (1, 2, 6) up to 2010 (red solid) and from 2011 on (black solid). Upper panel: Probability density, lower panel: Cumulative density.
[Colour figure can be viewed at wileyonlinelibrary.com]

intensive margin as the focus on such topics naturally crowds out other topics from a researcher's perspective. Also the attention of the discipline shifts if an increasing share of publications turns towards some specific topics.

Figure 7 looks at the density functions of the logarithmic sum of external topics as described above. One can observe the mean and the dispersion, that is, the second moment of the probability distribution. As Figure 6 highlights, there was a major surge in 2010. I split the distribution at this point, separating the time range in half and leading to $N_{\leq 2010} = 315$ and $N_{\leq 2010} = 462$. The earlier years' probability density function (PDF) is shaped close to a normally distributed PDF but with a fatter right tail. The PDF of the later years is bimodal with a second local maximum close to $log[\Sigma(\tau_1 + \tau_2 + \tau_6)] = 0$, which corresponds to a sum of probabilities ≈ 1 in levels.

Furthermore, one observes a shift to the right around the first local maximum. These characteristics of the PDFs correspond to a first-order stochastic dominance of the probability distribution of the earlier years over the later years, which can be drawn from the comparison of the cumulative density plots (CDF) in the lower panel of Figure 7. The corresponding one-sided Kolmogorov-Smirnov test confirms this as it rejects the null of no difference between the CDFs on every significance level (maximum difference $D^+ = 0.34,\ p = 0.000$). It is further empirical evidence for a shift toward this set of topics. Furthermore, the bimodality in recent years also hints at

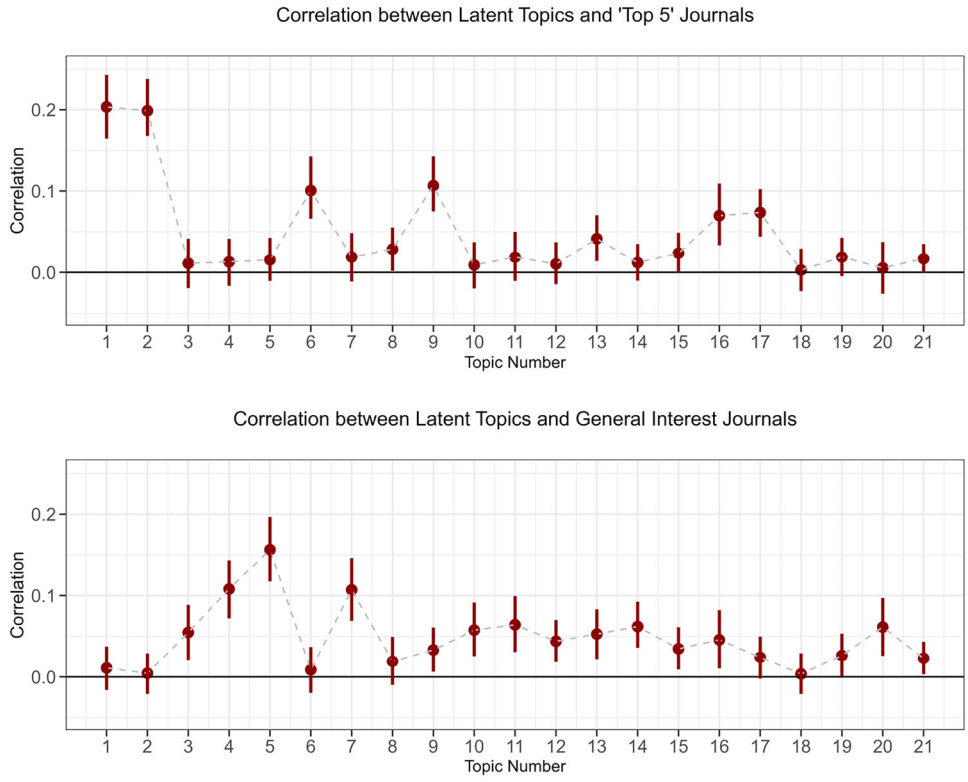

FIGURE 8 Topic correlation with journal types. Upper panel: 'Top 5' journals, lower panel: General interest journals. [Colour figure can be viewed at wileyonlinelibrary.com]

research "monocropping" in the sense that papers shift towards these topics and nearly solely focus on them, ignoring most of the other topics.

3.3 | Topic correlations with journal types

Until now, I have discussed topic prevalence without addressing where these topics are placed. In economics, however, the outlet of a paper has enormous importance (Fourcade et al., 2015; Schmal et al., 2023). This holds not only for a researcher's reputation but also for the audience it reaches. The so-called "Top 5" and the leading general interest journals are widely read and guide young researchers on where to look and what to investigate (Heckman & Moktan, 2020).

Figure 8 shows the correlation between the leading general interest journals and the 21 computed topics. The upper panel focuses on the "Top 5" journals, while the lower panel encompasses the leading "non-Top 5" general interest outlets, for example, the *Journal of the European Economic Association* or the *Economic Journal* of the Royal Economic Society. Eight topics are significantly correlated with the top five journals, namely 1, 2, 6, 8, 9, 13, 16, and 17. Meaningful are the extraordinarily high correlations with the first two topics, that is, those that take a rather general approach and those that discuss specific competition infringement cases. Furthermore, topics 6 and 9 are additional topics that address collusion cases and law enforcement on cartels

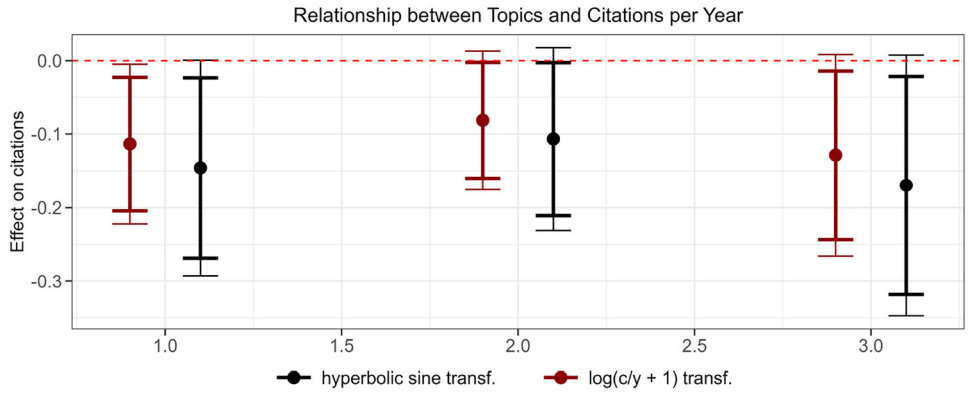

FIGURE 9 Relationship between topic prevalence and citations. Point estimates for the panel regression as described in Equation (1). Thick line: 90% confidence bands, thin line: 95% confidence band extensions. Left: topic 3, middle: topic 11, right: topic 19.
[Colour figure can be viewed at wileyonlinelibrary.com]

and anti-competitive institutions. Topics 16 and 17 tend to be case-related as well. Only topic 13 on vertical relations seems to diverge from this leading pattern.

Among the general interest journals shown in the lower panel of Figure 8, one finds lower correlations with many topics. The three most prominent connections exist with topics 4, 5, and 7, that is, on auction theory, information mechanisms, and principal-agent issues, which are also likely to be information-driven. Among field journals, dedicated IO journals, and antitrust journals, many topics are significantly correlated (as shown in Figure A.3 in the appendix), but there is not such a clear pattern as for general interest and Top 5 journals. The pattern for the latter is especially striking. The correlation with the first two topics is significant and much higher than any correlation with another topic. Together with topics 6, 9, and 17, these journals seem to focus intensely on past disclosed cartels. Other than topics of the categories "competition analysis" or "inside the cartel," those are rather case-study-based. see Figure A.4)

3.4 | Topic correlation with citations

Besides placing a paper in a highly-ranked journal, the number of citations of a paper is crucial in measuring academic impact and success. As my dataset allows me to look at the number of citations for each publication (up to early April 2022), I am able to study the relationship with the presence of certain topics. Conducting regression analyses for all 21 computed latent topics separately, I find significant relations for three topics. The results are shown in Figure 9. It shows two estimates for each topic. This stems from the fact that I report both transformations of the citations per year variable as described previously. All regression results, including those for $log(c/y)$, which suffer statistically from many missing values, are shown in the appendix in Table A.6.

The affected topics are those on market shares (3), capacity constraints (11), and procurement auctions (19). The latter two topics are part of the category "competition analysis" which is generally in decline as shown throughout this study. All three topics are negatively related to the number of annual citations. And all three are not significantly correlated with the "Top 5" journals. The other topics show no significant relation at all. Given the logarithmic nature of the

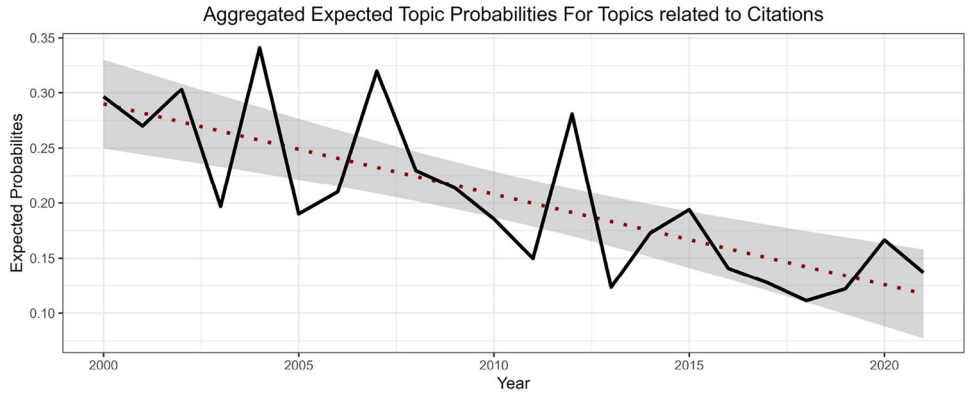

FIGURE 10 Mean topic prevalence of topics negatively correlated with citations. Black solid line: Topics correlated with citations (see Figure 9). Red dotted line: Fitted linear regression.
[Colour figure can be viewed at wileyonlinelibrary.com]

variables, a 1% increase in the expected probabilities of one of these topics is related to a decrease by around −0.1% up to −0.15%. One must remember that this is related to publications within the subdiscipline of collusion research.

As lower citation rates are an inferior outcome for researchers, this raises the question of the time dimension of the prevalence of these topics. Figure 10 shows that the share of topics with a lower level of citations diminished over time—even though it is fluctuating. Overall, researchers seem to realize these topics' relative unpopularity and shift their attention towards other topics. It is backed by the statistical tests for non-stationarity (see Table A.7 in the appendix). Again, the ADF test cannot reject the null of non-stationarity for two lags, while the PP test rejects it on the 5% level for types 2 and 3. As before, the KPSS test rejects its null of a unit root being absent, even on the 1% level. An alternative explanation might be that research on these particular topics has "ended" in the sense that there is (currently) nothing significantly new to contribute. Given that economics as a social science is never static and there is a continuous supply of new data, this seems less plausible.

There is sufficient evidence for non-stationarity, such that researchers tend to shift away from topics for which they can observe lower citations. It aligns with an NBER working paper studying European Research Council (ERC) grants. The authors find for rejected funding applicants a cutback in risky research after their repudiation (Veugelers et al., 2022). It is further evidence that researchers quickly adjust their agenda after receiving signals that their current approach does not pay off.

In contrast to the development of topics, the variety of subjects in papers has grown over time. I draw this from the neural network set up to understand similarities between papers year by year. I use the 'doc2vec' algorithm with 20 iterations of the training model and a 50×1 vector. As Mikolov et al. (2013) have shown, the dimension of the document vector matters. Figure A.6 in the appendix shows that variations (100×1, 500×1) do not affect the result. Figure 11 plots the change over time.

Unlike the previous analyses, the average similarity of papers *of a particular year* does not suggest an economically reasonable intertemporal relationship, suggesting a zero-lag specification. For a random walk without and with a drift (types 1 and 2), I cannot reject the null of

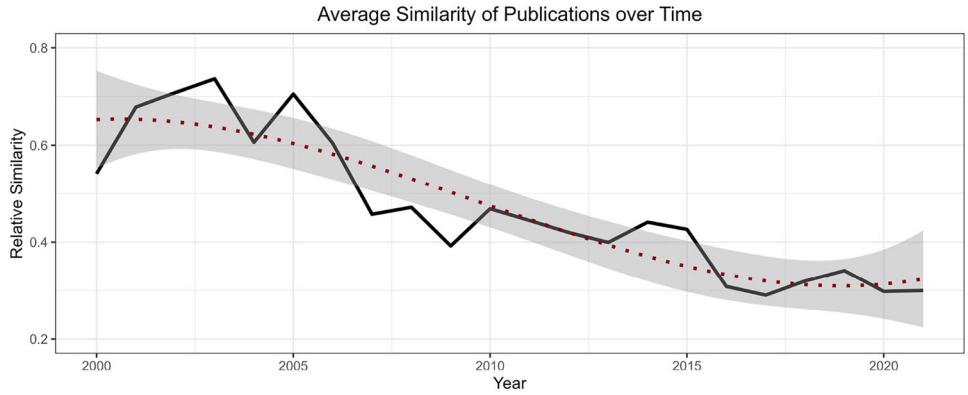

FIGURE 11 Development of content similarity of publications over time. Black solid line: Mean similarity of papers of a specific year with all papers of the same year excluding a paper's similarity with itself. Red dotted line: Fitted linear polynomial regression of order 3. Neural network specification: paragraph vector dimension: 50. Training iterations: 20. Method: paragraph vector with distributed memory (PV-DM, see Le & Mikolov, 2014, for details).
[Colour figure can be viewed at wileyonlinelibrary.com]

non-stationarity using ADF and PP. The KPSS test also suggests the presence of non-stationarity. The results are shown in the appendix in Table A.8.

As I find a notable decline in similarity measured with a plain univariate linear regression of relative similarity on time (coefficient: −0.0199, $p = .0016$), it implies that the variety of subjects actually increases. In particular, similarity within the body of research diminishes by some 2pp per year, which leads to a halving of the overall proximity of contents over time. Thus, the antitrust community faces two parallel developments: We see growing diversity in the subjects of papers, for example, different industries in which firms collude. However, at the same time, there is evidence for a sharp narrowing in underlying topics. Put differently, the scope of research diminishes while the scope of applications grows.[21]

3.5 | Reciprocal topic prevalence over time

While the preceding analyses already bear essential insights into how collusion research evolved during the early 21st century, a multivariate examination may provide a further understanding of the co-movement of topics discussed. Particularly meaningful is how the rise and fall of topic prevalence affect other topics. Causal inference in a microeconometric sense is difficult to obtain given the lack of a major shock that could serve as causal identification. There is some discontinuity for case study topics in 2009 (see Figure 6). The global financial crisis coincides with it. However, there is no apparent reason why this should have affected research on collusion, especially in light of the submission-publication lag. Therefore, I utilize the toolbox of macroeconomic multivariate time series analysis instead. Here, I look at the joint evolution of the "case-study topics" (topics 1, 2, and 6) and the "competition analysis" category (topics 4, 7, 11, 12, 18, 19).

To do that, I set up a Vector Autoregressive (VAR) model.[22] Even though there is evidence for cointegration, I utilize the Toda-Yamamoto (1995) approach (TY), which states that a VAR in base numbers instead of differences can be fully used even if there exists cointegration as long as one

TABLE 3 Granger causality wald tests.

Dependent variable	Excluded	χ^2	df	Prob. $> \chi^2$
External topics	competition analysis	21.653	3	.000***
Competition analysis	external topics	2.729	3	.435

Signif. codes: ***: $p < .01$, **: $p < .05$, *: $p < .1$.

adds the order of integration to the lag length of the VAR and the order of cointegration does not exceed the initial leg length. The selection criteria suggest two lags and cointegration of order 1 (see Tables A.9 and A.11 in the appendix). This leads to a specification with three lags (see Table A.10)

The technical evaluation shows that this VAR (using logarithmic values) satisfies the common criteria for a reliable model (see Tables A.12 and A.13 in the appendix). All eigenvalues are inside the unit circle and satisfy the stability condition of VAR models (see Table A.14 in the appendix). The residuals are not auto-correlated and satisfy all conditions for a well-behaved normal distribution. A convenient feature of VAR models is the Granger causality analysis. Table 3 shows the results for the very compact Granger causality Wald tests within the two-variable setting. Essentially, they test whether the past values of a variable influence another variable: If so, the excluded variable "Granger causes" the other one. Thus, it only satisfies a weak definition of causality.[23]

Past values of the "competition analysis" influence the prevalence of the "external topics" while the opposite does not hold because the test statistic cannot reject the null of no effect. It is reflected in the impulse response functions presented in the upper panels of Figure 12. One observes that a shock to the external topics (LHS) does not affect the "competition analysis." In contrast, we observe a significant decrease for two periods in the external topics following a positive shock to the competition-related topics, as the right panel in the figure shows.

The lower panels of Figure 12 show the forecast error variance decomposition (FEVD), which adds additional evidence for the previous finding. It measures the share a shock happening to one variable has on the forecast error of the other. The lower left panel shows that a shock to the external variables has no significant effect on the forecast error of the competition analysis. Hence, the latter evolves independently of the exogenous topics. The opposite is true for the right panel. It takes one period of inertia, but from the second lead on, a shock of the competition analysis category significantly and persistently affects the evolution of the external topic prevalence and effectively predicts a portion of slightly below 1/2 of it (see Table A.15 in the appendix for details). Hence, the external topics are at least partially endogenous because their evolution is driven by the changes in the prevalence of competition analysis in the literature but not vice versa.

4 | CONTEXTUALIZATION AND CONCLUSION

This systematic review of 777 academic publications on collusion identifies essential insights that may affect antitrust authorities and practitioners in their ability to detect cartels and collusive behavior. Using a structural topic model, I am able to identify the underlying latent topics that have driven collusion research in the past two decades. I obtain a hitherto non-existent research clustering in this subdiscipline by doing so. I have conducted various univariate time series analyses based on the computed topic probabilities. They reveal a substantial decline in competition analysis, often based on game-theoretical considerations and so-called toy models that study a particular feature of a market and its partial equilibrium effects.

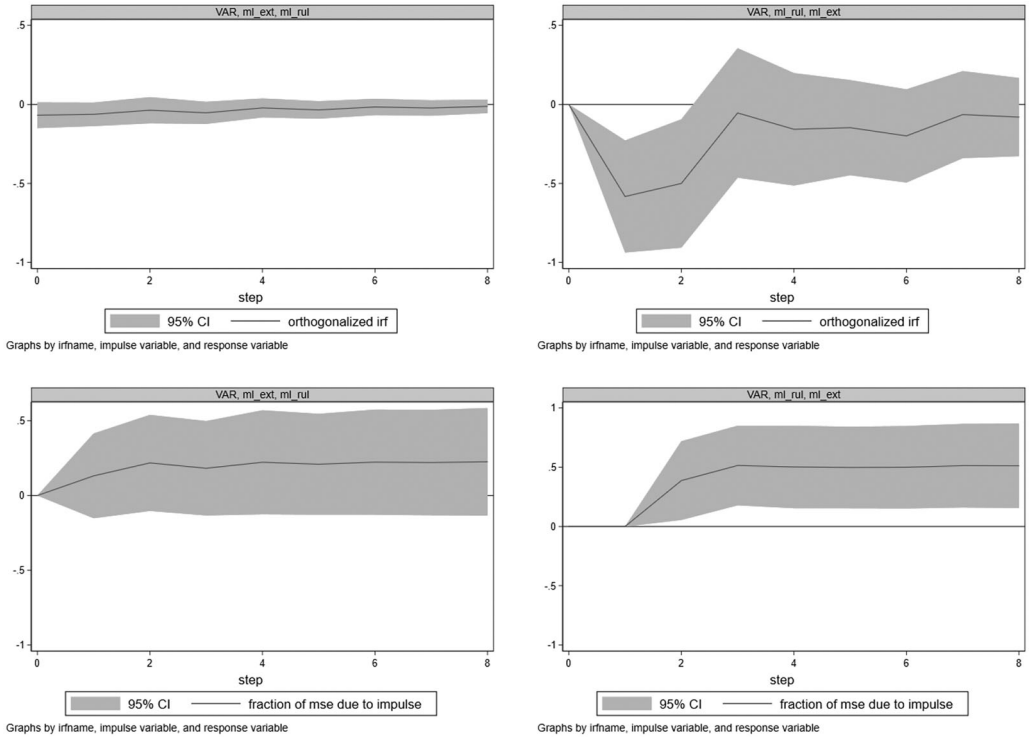

FIGURE 12 Impulse response functions and forecast error variance decomposition. LHS: Impulse variable: "External topics," response variable: "competition analysis" RHS: vice versa. Upper panels: Impulse response functions, lower panels: forecast error variance decomposition.

Further, I have computed panel regressions investigating the relation between latent topics and citation rates. While there are only a few significant correlations, all of them suggest a negative relationship. As shown in the time series plot on the expected mean topic prevalence, one sees a downward trend for the topics with a negative relation to citations. Hence, researchers select these topics, maybe noting that the receptions (captured by citations) are low in the discipline and might adjust their research foci over time. While researchers tend to narrow the number of topics addressed in their papers, they get more diverse as a neural network-based analysis on the similarity of the body of research can show.

Last, I have conducted a multivariate analysis. It reveals that the decline in game theory is not triggered by the new dominance of data-heavy case studies but instead fills the gap the turn away from competition analysis has left. All of this happened parallel to a steady broadening of contents. While the core issues of the papers seem to be shrinking, the overall contents, such as cases discussed, grow. Hence, we have a two-level decoupling in research: Fewer topics get wider attention and coverage than two decades ago.

It is a positive development that the quantitative rigor of these case studies has increased a lot since the "credibility revolution" (Angrist & Pischke, 2010) in empirical economics. For example, causal inference designs are nowadays widespread in such papers. Furthermore, structural approaches in empirical industrial organization may be helpful to avoid the reliance on ad hoc assumptions.

The downside is that case studies, by definition, are backward-looking. Whether this will harm cartel detection remains an open question, as rule-based approaches may be more adoptable by competition authorities than granular but case-specific analyses of past cartels.[24] It is notable, though, that case study topics have by far the highest correlation with the "Top 5" journals in economics. While one-third of the topics are significantly correlated with these prestigious outlets, the first two topics have the highest point estimates for correlation across all topics and journal categories.

It implies that the leading figures of the economic discipline particularly value this kind of paper. Naturally, it should foster many follow-up papers of IO researchers on various cartels. Hence, it is likely that these cartel case studies remain a strong pillar of collusion research. Taking a closer look at IO field journals such as the *Journal of Industrial Economics* or the *International Journal of Industrial Organization*, which publish a lot of collusion research, it turns out that they still focus on competition analysis and market outcomes in terms of topics.

The "game-theoretic oligopoly theory" (Budzinski, 2007, p. 301) pushed forward in the 1990s and considered a threat to competition economics and antitrust policy appears to be in retreat in collusion research. From the availability of large datasets and computing capacity emerge sophisticated and novel contributions—as predicted by Kovacic and Shapiro (2000). In contrast, the old Harvard school, predominant from the 1930s on and based mainly on industry case studies (Bresnahan & Schmalensee, 1987), seems to experience a revival. The high demand for validations of the core findings may require such detailed data that can be only drawn from case studies. When reviewing the large mobile-phone spectrum auctions at the turn of the millennium, Klemperer (2002) noted that it mostly needs "elementary economics" and what is usually understood as rule and incentive-based thinking, namely the avoidance of collusion and misbehavior of the auction participants.

Already in 2007, Caves highlighted the benefits of the "old IO" approach based on cross-sectional analyses. While I am far away from criticizing or questioning quantitative rigor, the developments in the IO literature outlined in this paper put in question to which extent competition authorities get equipped with the right tools to investigate markets and detect collusion.[25] This paper also makes no claims on the usefulness of sophisticated empirical models such as structural estimations. Nevertheless, in the field of collusion, Ghosal and Sokol (2014) have shown that cartel enforcement in the US has shifted from a large number of cartels detected in the 1980s and 1990s towards an approach that detects a comparatively low number of cartels. Among them are very large cartels, and high fines were imposed. This shift from many (small) cartels towards a few "big fishes" might correspond to the shift in the literature.

In the very last step, I want to provide some avenues for potential policy interventions. An important way of steering research foci is the creation of incentives. Hottenrott and Lawson (2014) and Schmal (2023) could show that external funding may affect researcher behavior. Hence, research organizations and competition authorities could foster research on topics that have lost attention by academics by handing out grants and research support. Furthermore, authorities could host or support academic conferences that address specific topics. It would raise attention to these topics and bring together dispersed researchers who might start new collaborations. More extensive approaches would be funding chairs or research centers at universities. While it is a rather costly exercise, hosting conferences would be a low-threshold measure with comparatively low costs (for a public body), which may still have the potential to evoke a sizable echo in the academic community, especially if such conferences become regular annual events.

ACKNOWLEDGMENT

I am very grateful for valuable feedback from the handling editor Les Oxley, two anonymous referees, Oliver Budzinski, Justus Haucap, Paul Hünermund, Ekkehard Köhler, Anselm Küsters, Maikel Pellens, and Chris Snyder, as well as feedback from participants at the 16th Biennal Conference of the International Network for Economic Method, the Annual Conference of the Royal Economic Society 2023, the MSI Research Group Meeting at KU Leuven, the DICE PhD Research Workshop, the 2022 NOUS Workshop at Albert Ludwigs University Freiburg, and the 56th Hohenheim colloquium. All errors are my own. Funding by the German Research Foundation (DFG) is gratefully acknowledged (Funding No. #235577387/GRK 1974).

CONFLICT OF INTEREST STATEMENT

The author declares that there exist no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The writing, spelling, and grammar of this research paper have been enhanced through the use of artificial intelligence (AI) tools. The AI technology was employed to aid in identifying and correcting errors in writing, spelling, and grammar, while human judgment was exercised in reviewing and approving the AI-generated suggestions, ensuring that the content accurately conveyed the intended meaning, and aligning it with the research findings.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from Scopus. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors with the permission of Scopus. All code files as well as all data computed by the author are available on GitHub.

ORCID

W. Benedikt Schmal https://orcid.org/0000-0003-2400-2468

ENDNOTES

[1] See https://competition-policy.ec.europa.eu/document/download/b19175c3-c693-410b-b669-27d4360d359c_en?filename=cartels_cases_statistics.pdf(up to November 2022).

[2] The exchange between academics and practitioners has been particularly established by the Chicago school of antitrust in the United States and the "more economic approach" for competition decisions in the EU.

[3] Using JEL codes to identify content-wise related papers would be an alternative, but antitrust journals often do not use them. Furthermore, researchers could be biased or strategically act when assigning specific codes to their papers (Wehrheim, 2019). Keywords would be another option, but latent topics go into much more detail and circumvent the subjective self-reporting by the authors. A different approach would be gathering a corpus of economics and antitrust journals, running a topic modeling on all papers, and extracting those with a prevalence of topics related to collusion. In the second step, a topic model could be constructed for this subset of publications. The core difference would be that the papers are not selected by a manual keyword identification. It would prevent papers from being falsely ignored if they discuss collusion without using the mentioned terms. While this is an advantage, the two-step process still requires a manual selection of topics that one may consider as collusion-related. Furthermore, especially in the abstracts, authors need to use common keywords to inform potential readers about the paper. Thus, abstracts are unlikely to miss out on the core keywords of the strand of literature they belong to. Thus, authors are unlikely to use uncommon and creative language in this part of a research article.

[4] I technically also exclude those papers without any abstract listed in Scopus. Table A.1 in the appendix lists the journals.

[5] Of course, analyzing full papers would provide much more granularity regardless of the occurrence of intruder terms. However, this paper takes the approach of investigating major developments in the literature. As abstracts serve as condensed summaries of what a paper is about, they should be a strong approximation of the content of the full papers. Furthermore, while papers changed a lot throughout the last decades, especially with the emergence of extensive online appendices and empirical robustness checks, abstracts remain widely unchanged as their purpose remains unchanged regardless of whether the paper is distributed online or in print. Thus, abstracts are preferable as their size arguably did not change in the past two decades while the size of journal articles increased a lot. The top five economics journals faced, for example, an increase by a factor of four between 1970 and 2010 (Card & DellaVigna, 2014). Therefore, abstracts continue to provide all essential information on the topics covered by the paper and are not distorted by the length inflation of full articles.

[6] I use the widely applied list created and published by Lewis et al. (2004). The list is publicly available at MIT's Computer Science and Artificial Intelligence Laboratory: http://www.ai.mit.edu/projects/jmlr/papers/volume5/lewis04a/a11-smart-stop-list/english.stop

[7] Table A.3 in the appendix lists these additional stopwords.

[8] I use the letter K to choose the optimal number of topics as done in the literature. However, later on, I label topics with the Greek letter τ_k, which abbreviates topic $k \leq K$.

[9] Consider as an example the "document" d_1 ="Hello world", that contains of w_1 = hello and w_2 = world and, say, one topic k_1 = greeting.

[10] These categories are "top 5" journals, general interest (without top 5), microeconomic field journals, specific IO journals, and antitrust journals. The so-called top five journals in economics are *the American Economic Review*, *Econometrica*, *the Journal of Political Economy*, *the Quarterly Journal of Economics*, and *the Review of Economic Studies* (ordered alphabetically). They have stellar importance within the discipline. While IO journals nearly exclusively focus on economic elaborations of (in my case) collusion, antitrust journals also encompass legal and practitioners' perspectives.

[11] See, for example, Hadavand et al. (ming), Björk and Solomon (2013).

[12] This is backed by the lag heuristic as part of the computation approach for a heteroskedasticity and autocorrelation consistent covariance matrix as suggested by Newey and West (1994). Given the structure of the data, we need $n \in \mathbb{N}$, such that I set $n^* \equiv \lfloor n \rfloor$ as done by Schwert (1989).

[13] The use of the corresponding author makes the regression more convenient than using a factor variable for each group of authors. In 131 papers, no corresponding authors are mentioned, especially in older publications. In these cases, I replace the missing values with the author teams.

[14] Analyzing single topics would be interesting. Given the comparatively large amount of topics and the limited length of the time series, categories as clustered in Table 2 are more suitable.

[15] For example, Küsters (2022) draws in his analysis from a set of "lessons from the past" in monetary policy to cluster topics. Wehrheim (2021) uses a combination of the cosine similarity of topics (Aletras & Stevenson, 2014) and human reasoning to group topics to superior meta-topics. It bears the problem of technical similarity but not necessarily content similarity. Wehrheim (2019) compares the estimations from different K specifications to decide which one best serves the paper's purpose. In general, iterative experimentation and human reasoning are broadly applied when it comes to the decision of how many topics one should compute, as technical computations of the optimal number may only provide a technical answer to this question, which might be insufficient when it comes to interpreting the topic content and further analyzing it.

[16] As a robustness check, an alternative specification with $\hat{K} = 10$ topics is shown in Table B.1 in Appendix B. It shows a clustering with many similarities to the one chosen, which strengthens the reliability of the setting with $K = 21$. However, it misses many details due to the low number of topics.

[17] See for an introduction to leniency programs, for example, Hinloopen (2003).

[18] As a robustness check, I hand-code the predominant methodology of a paper, that is, whether it is mainly theoretical, empirical, or experimental in its research design. Figure A.2 in the appendix plots the evolution of methodological types over time.

[19] Chow tests for a structural break in 2009, 2010, and 2011 cannot reject the null hypothesis of no break.

[20] See, for example, the EU Summary Report on Cartel Statistics: https://ec.europa.eu/competition-policy/system/files/2022-07/cartels_cases_statistics.pdf

[21] Of course, one has to be aware that this is a semantic analysis. Hence, it could be the case that abstracts become more diverse in terms of the wording used without a higher diversity in actual content. Some suggestive evidence

for this is provided by Wang et al. (2022) as they detect a decrease in the readability of abstracts over time, which implies a more sophisticated sentence structure, which might be only based on more complicated grammar but could also be related to a higher usage of rare words. Nevertheless, it does not invalidate my initial hypothesis that research on collusion—while it becomes more diverse in some dimensions—the actual variety of topics is ultimately in decline.

[22] I test for cointegration between the variables, that is, a long-run relationship between some or all of them. The Johansen tests suggest cointegration of order 2. The Engle-Granger test finds evidence for cointegration at all (the order is not predictable here); see Table A.9 in the appendix for details).

[23] See for an elaboration on the causality dimension, for example, Granger (1988) or Maziarz (2015).

[24] In competition policy more broadly, however, rule-based thinking is again on the rise, considering recent legislative EU reforms such as the digital markets act (see for details, e.g., Larouche & de Streel, 2021). I thank one of the anonymous reviewers for particularly pointing at it.

[25] It has to be acknowledged, however, that competition authorities nowadays also apply topical computational and algorithmic tools to detect cartels automatically (see, e.g., Harrington & Imhof, 2022; Sanchez-Graells, 2019). While limited availability of data on cartels might be a challenge to these approaches (de Cooman, 2023), they could be beneficial in screening publicly observable actions such as bidding behavior (Huber & Imhof, 2019; Imhof & Wallimann, 2021).

REFERENCES

Airoldi, E. M., & Bischof, J. M. (2016). Improving and evaluating topic models and other models of text. *Journal of the American Statistical Association*, *111*(516), 1381–1403.

Aletras, N., & Stevenson, M. (2014). Measuring the similarity between automatically generated topics. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, vol. 2: Short Papers*, (pp. 22–27).

Amari, S.-I. (1995). Information geometry of the em and em algorithms for neural networks. *Neural Networks*, *8*(9), 1379–1408.

Ambrosino, A., Cedrini, M., Davis, J. B., Fiori, S., Guerzoni, M., & Nuccio, M. (2018). What topic modeling could reveal about the evolution of economics. *Journal of Economic Methodology*, *25*(4), 329–348.

Angrist, J. D., & Pischke, J.-S. (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives*, *24*(2), 3–30.

Ash, E., Chen, D. L., & Naidu, S. (2022). Ideas have consequences: The impact of law and economics on american justice. Technical report, NBER Working Paper 29788.

Bischof, J., & Airoldi, E. M. (2012). Summarizing topical content with word frequency and exclusivity. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, (pp. 201–208).

Björk, B.-C., & Solomon, D. (2013). The publishing delay in scholarly peer-reviewed journals. *Journal of Informetrics*, *7*(4), 914–923.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, *3*(Jan), 993–1022.

Bolotova, Y. V. (2009). Cartel overcharges: An empirical analysis. *Journal of Economic Behavior & Organization*, *70*(1-2), 321–341.

Boyer, M., & Kotchoni, R. (2015). How much do cartel overcharge? *Review of Industrial Organization*, *47*(2), 119–153.

Bresnahan, T., & Schmalensee, R. (1987). The empirical renaissance in industrial economics: An overview. *Journal of Industrial Economics*, 371–378.

Budzinski, O. (2007). Monoculture versus diversity in competition economics. *Cambridge Journal of Economics*, *32*(2), 295–324.

Card, D., & DellaVigna, S. (2014). Page limits on economics articles: Evidence from two journals. *Journal of Economic Perspectives*, *28*(3), 149–168.

Caves, R. E. (2007). In praise of the old IO. *International Journal of Industrial Organization*, *25*(1), 1–12.

Cherrier, B., & Svorenčík, A. (2018). The quantitative turn in the history of economics: Promises, perils and challenges. *Journal of Economic Methodology*, *25*(4), 367–377.

Connor, J. M. (2008). *Global price fixing* (2nd updated and revised edition ed.). Volume 26 of *Studies in industrial organization*. Springer.

Davidson, R., & MacKinnon, J. G. (1993). *Estimation and inference in econometrics* (1st ed.). Oxford University Press.
de Cooman, J. (2023). Outsmarting pac-man with artificial intelligence, or why AI-driven cartel screening is not a silver bullet. *Journal of European Competition Law & Practice*, *14*(4), 186–202.
Ebadi, A., Xi, P., Tremblay, S., Spencer, B., Pall, R., & Wong, A. (2021). Understanding the temporal evolution of COVID-19 research through machine learning and natural language processing. *Scientometrics*, *126*(1), 725–739.
Einav, L., & Levin, J. (2010). Empirical industrial organization: A progress report. *The Journal of Economic Perspectives*, *24*(2), 145–162.
Ezrachi, A., & Stucke, M. E. (2016). Virtual competition. *Journal of European Competition Law & Practice*, *7*(9), 585–586.
Ferrara, F. M., Masciandaro, D., Moschella, M., & Romelli, D. (2022). Political voice on monetary policy: Evidence from the parliamentary hearings of the European Central Bank. *European Journal of Political Economy*, *74*, 102143.
Fourcade, M., Ollion, E., & Algan, Y. (2015). The superiority of economists. *Journal of Economic Perspectives*, *29*(1), 89–114.
Gaule, P., & Maystre, N. (2011). Getting cited: Does open access help? *Research Policy*, *40*(10), 1332–1338.
Gentzkow, M., Kelly, B., & Taddy, M. (2019). Text as data. *Journal of Economic Literature*, *57*(3), 535–574.
Ghosal, V., & Sokol, D. D. (2014). The evolution of us cartel enforcement. *The Journal of Law and Economics*, *57*(S3), S51–S65.
Granger, C. W. (1988). Some recent development in a concept of causality. *Journal of Econometrics*, *39*(1-2), 199–211.
Green, E. J., Marshall, R. C., & Marx, L. M. (2014). Tacit collusion in oligopoly. *The Oxford handbook of international antitrust economics* (Vol. 2, pp. 464–497).
Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, *21*(3), 267–297.
Gropp, C., Herzog, A., Safro, I., Wilson, P. W., & Apon, A. W. (2016). Scalable dynamic topic modeling with clustered latent Dirichlet allocation (CLDA). arXiv Preprint, arXiv:1610.07703.
Hadavand, A., Hamermesh, D. S., & Wilson, W. W. (forthcoming). Publishing economics: How slow? why slow? is slow productive? how to fix slow? *Journal of Economic Literature*.
Hansen, S., & McMahon, M. (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics*, *99*, S114–S133.
Hansen, S., McMahon, M., & Prat, A. (2018). Transparency and deliberation within the FOMC: A computational linguistics approach. *The Quarterly Journal of Economics*, *133*(2), 801–870.
Hansen, S., McMahon, M., & Tong, M. (2019). The long-run information effect of central bank communication. *Journal of Monetary Economics*, *108*, 185–202.
Harrington, J., & Imhof, D. (2022). Cartel screening and machine learning. *Stanford Computational Antitrust*, *2*, 133–154.
Heckman, J. J., & Moktan, S. (2020). Publishing and promotion in economics: The tyranny of the top five. *Journal of Economic Literature*, *58*(2), 419–470.
Hinloopen, J. (2003). An economic analysis of leniency programs in antitrust law. *De Economist*, *151*(4), 415–432.
Hottenrott, H., & Lawson, C. (2014). Research grants, sources of ideas and the effects on academic research. *Economics of Innovation and New Technology*, *23*(2), 109–133.
Hovenkamp, H. (2018). Whatever did happen to the antitrust movement. *Notre Dame Law Review*, *94*, 583–638.
Hovenkamp, H. (2021). The looming crisis in antitrust economics. *Boston University Law Review*, *101*, 489–545.
Huber, M., & Imhof, D. (2019). Machine learning with screens for detecting bid-rigging cartels. *International Journal of Industrial Organization*, *65*, 277–301.
Imhof, D., & Wallimann, H. (2021). Detecting bid-rigging coalitions in different countries and auction formats. *International Review of Law and Economics*, *68*, 106016.
Jones, A., & Kovacic, W. E. (2020). Antitrust's implementation blind side: challenges to major expansion of us competition policy. *The Antitrust Bulletin*, *65*(2), 227–255.
Klemperer, P. (2002). What really matters in auction design. *Journal of Economic Perspectives*, *16*(1), 169–189.
Kovacic, W. E., & Shapiro, C. (2000). Antitrust policy: A century of economic and legal thinking. *Journal of Economic Perspectives*, *14*(1), 43–60.

Kumar, V., & Srivastava, A. (2022). Trends in the thematic landscape of corporate social responsibility research: A structural topic modeling approach. *Journal of Business Research*, *150*, 26–37.

Küsters, A. (2022). Applying lessons from the past? Exploring historical analogies in ECB speeches through text mining, 1997–2019. *International Journal of Central Banking*, *18*(1), 277–329.

Kwiatkowski, D., Phillips, P. C., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, *54*(1-3), 159–178.

Lämmel, R., & Jones, S. P. (2003). Scrap your boilerplate: A practical design pattern for generic programming. *Association for Computing Machinery SIGPLAN Notices*, *38*(3), 26–37.

Lancieri, F., Posner, E. A., & Zingales, L. (2022). The political economy of the decline of antitrust enforcement in the united states. Working Paper 30326, NBER.

Larouche, P., & de Streel, A. (2021). The european digital markets act: A revolution grounded on traditions. *Journal of European Competition Law & Practice*, *12*(7), 542–560.

Larsen, V. H. (2021). Components of uncertainty. *International Economic Review*, *62*(2), 769–788.

Le, Q., & Mikolov, T. (2014). Distributed representations of sentences and documents. In *International Conference on Machine Learning*, (pp. 1188–1196). PMLR.

Levenstein, M. C., & Suslow, V. Y. (2006). What determines cartel success? *Journal of Economic Literature*, *44*(1), 43–95.

Lewis, D. D., Yang, Y., Russell-Rose, T., & Li, F. (2004). Rcv1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, *5*(Apr), 361–397.

Leybourne, S. J., & Newbold, P. (1999). The behaviour of Dickey–Fuller and Phillips-Perron tests under the alternative hypothesis. *The Econometrics Journal*, *2*(1), 92–106.

Maziarz, M. (2015). A review of the granger-causality fallacy. *The Journal of Philosophical Economics*, *8*(2), 86–105.

McCabe, M., & Snyder, C. (2014). Identifying the effect of open access on citations using a panel of science journals. *Economic Inquiry*, *52*(4), 1284–1300.

McCabe, M., & Snyder, C. (2021). Cite unseen: Theory and evidence on the effect of open access on cites to academic articles across the quality spectrum. *Managerial and Decision Economics*, *42*(8), 1960–1979.

McGinnis, M. D. (2011). An introduction to IAD and the language of the Ostrom workshop: A simple guide to a complex framework. *Policy Studies Journal*, *39*(1), 169–183.

Mehra, S. K. (2015). Antitrust and the robo-seller: Competition in the time of algorithms. *Minnesota Law Review*, *100*, 1323.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arχiv preprint arχiv:1301.3781*.

Mimno, D., Wallach, H., Talley, E., Leenders, M., & McCallum, A. (2011). Optimizing semantic coherence in topic models. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, (pp. 262–272).

Moschella, M., Pinto, L., & Martocchia Diodati, N. (2020). Let's speak more? How the ECB responds to public contestation. *Journal of European Public Policy*, *27*(3), 400–418.

Nathan, M., & Rosso, A. (2022). Innovative events: Product launches, innovation and firm performance. *Research Policy*, *51*(1), 104373.

Newey, W. K., & West, K. D. (1994). Automatic lag selection in covariance matrix estimation. *Review of Economic Studies*, *61*(4), 631–653.

OECD. (2017). Algorithms and collusion: Competition policy in the digital age. Published September 14, 2017, URL: https://www.oecd.org/competition/algorithms-collusion-competition-policy-in-the-digital-age.htm

Ostrom, E. (2005). Understanding the diversity of structured human interactions. In *Understanding institutional diversity* (pp. 3–29). Princeton University Press.

Ostrom, E. (2011). Background on the institutional analysis and development framework. *Policy Studies Journal*, *39*(1), 7–27.

Potters, J., & Suetens, S. (2013). Oligopoly experiments in the current millennium. *Journal of Economic Surveys*, *27*(3), 439–460.

Qasim, M. (2017). Sustainability and wellbeing: A scientometric and bibliometric review of the literature. *Journal of Economic Surveys*, *31*(4), 1035–1061.

Roberts, M. E., Stewart, B. M., & Tingley, D. (2019). Stm: An r package for structural topic models. *Journal of Statistical Software*, *91*, 1–40.

Roberts, M. E., Stewart, B. M., Tingley, D., & Airoldi, E. M. (2013). The structural topic model and applied social science. In *Advances in neural information processing systems workshop on topic models: Computation, application, and evaluation* (Vol. 4, pp. 1–20). Harrahs and Harveys, Lake Tahoe.

Rose, M. E., & Kitchin, J. R. (2019). pybliometrics: Scriptable bibliometrics using a Python interface to Scopus. *SoftwareX*, *10*, 100263.

Sanchez-Graells, A. (2019). 'Screening for cartels' in public procurement: Cheating at solitaire to sell fool's gold? *Journal of European Competition Law & Practice*, *10*(4), 199–211.

Schinkel, M. P., & Treuren, L. (2021). Green antitrust:(more) friendly fire in the fight against climate change. *Available at SSRN #3749147*, pp. 2020–72.

Schmal, W. B. (2022a). Governing the cartels – An institutional exploration. Available at SSRN:https://dx.doi.org/10.2139/ssrn.4430726

Schmal, W. B. (2022b). Legislative flagships: Spillovers of law amendments to economic research in associated jurisdictions. unpublished manuscript.

Schmal, W. B. (2023). The X factor: Open access, new journals, and incumbent competitors. MSI working paper 2307, KU Leuven, Faculty of Economics and Business (FEB), Department of Management, Strategy, and Innovation (MSI).

Schmal, W. B., Haucap, J., & Knoke, L. (2023). The role of gender and coauthors in academic publication behavior. *Research Policy*, *52*(10).

Schwert, G. W. (1989). Tests for unit roots: A Monte Carlo investigation. *Journal of Business & Economic Statistics*, *7*(2), 5–17.

Shapiro, C. (2020). Antitrust: What went wrong and how to fix it. *Antitrust*, *35*, 33–45.

Sybrandt, J., Carrabba, A., Herzog, A., & Safro, I. (2018). Are abstracts enough for hypothesis generation? In *2018 IEEE International Conference on Big Data* (pp. 1504–1513). IEEE.

Toda, H. Y., & Yamamoto, T. (1995). Statistical inference in vector autoregressions with possibly integrated processes. *Journal of Econometrics*, *66*(1-2), 225–250.

Veugelers, R., Wang, J., & Stephan, P. (2022). Do funding agencies select and enable risky research: Evidence from erc using novelty as a proxy of risk taking. Working Paper 30320, NBER.

Wang, S., Liu, X., & Zhou, J. (2022). Readability is decreasing in language and linguistics. *Scientometrics*, *127*(8), 4697–4729.

Wehrheim, L. (2019). Economic history goes digital: Topic modeling the journal of economic history. *Cliometrica*, *13*(1), 83–125.

Wehrheim, L. (2021). *Im Olymp der Ökonomen: zur öffentlichen Resonanz wirtschaftspolitischer Experten von 1965 bis 2015*. Mohr Siebeck.

Wils, W. P. (2007). The european commission's 2006 guidelines on antitrust fines: A legal and economic analysis. *World Competition*, *30*(2).

Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature*, *566*, 378–382.

How to cite this article: Schmal, W. B. (2024). Vice versa: The decoupling of content and topic heterogeneity in collusion research. *Journal of Economic Surveys*, *38,* 1686–1730. https://doi.org/10.1111/joes.12600

APPENDIX A

TABLE A.1 List of journals with the number of included articles.

Journal	#Articles	Share	Cum.	*Type*
International Journal of Industrial Organization	134	17.25	17.25	IO
Review of Industrial Organization	81	10.42	27.67	IO
Journal of Economic Behavior and Organization	64	8.24	35.91	IO
RAND Journal of Economics	49	6.31	42.21	IO
Journal of Industrial Economics	48	6.18	48.39	IO
Antitrust Bulletin	46	5.92	54.31	AT
Journal of Economic Theory	39	5.02	59.33	FI
Journal of Antitrust Enforcement	35	4.50	63.84	AT
European Economic Review	29	3.73	67.57	GI
World Competition	27	3.47	71.04	AT
European Competition Journal	23	2.96	74.00	AT
Journal of Law and Economics	22	2.83	76.83	FI
Management Science	22	2.83	79.67	FI
Journal of Economics and Management Strategy	20	2.57	82.24	FI
American Economic Review	15	1.93	84.17	T5
International Economic Review	13	1.67	85.84	GI
Review of Economic Studies	13	1.67	87.52	T5
Econometrica	12	1.54	89.06	T5
Economic Journal	12	1.54	90.60	GI
Journal of Law, Economics, and Organization	12	1.54	92.15	FI
Review of Economics and Statistics	12	1.54	93.69	GI
American Economic Journal: Microeconomics	8	1.03	94.72	FI
Journal of International Economic Law	8	1.03	95.75	AT
Journal of Political Economy	8	1.03	96.78	T5
Journal of the European Economics Association	7	0.90	97.68	GI
Journal of European Comp. Law and Practice	5	0.64	98.33	AT
Theoretical Economics	4	0.51	98.84	FI
American Economic Journal: Applied Econ.	2	0.26	99.10	GI
Economic Policy	2	0.26	99.36	GI
Journal of Economic Literature	2	0.26	99.61	GI
Quarterly Journal of Economics	2	0.26	99.87	T5
Journal of Economic Perspectives	1	0.13	100	GI
Total	**777**			

Journals ordered by the number of articles published in the branch of collusion. Journal types (manually assigned): AT – Antitrust; FI – Field; GI – General Interest (w/o Top 5); IO – Industrial Organization; T5 – Top 5 Journals.

TABLE A.2 Paper distribution by journal type and publication time.

	2000 - 04	2005 - 09	2010 - 14	2015 - 19	≥ 2020	Total
Antitrust	1	3	33	72	35	144
Field	22	24	28	38	15	127
General Int.	17	19	19	19	6	80
Ind. Org.	69	105	65	92	45	376
Top 5 Journals	15	12	7	11	5	50
Total	124	163	152	232	106	**777**

TABLE A.3 Custom stopwords additional to the 'SMART' stopwords (Lewis et al., 2004).

(1)	(2)	(3)	(4)	(5)
blackwell	journal	academic	editorial	cartel
rand	institute	among	paper	cartels
springer	ltd	supervisor	section	collusion
north-holland	llc	can	literature	cartels
sciencebusiness	iii	due	discuss	collusive
elsevier	press		may	
	eu		show	
			mine	
			will	

TABLE A.4 Unit root tests for non-stationarity in 'competition analysis' topics.

Non-stationarity type	*lag*	*statistic*	*p-value*
Augmented Dickey-Fuller Test			
type 1	0	−1.4285	.1607
type 1	1	−0.7044	.4170
type 1	2	−1.4713	.1456
type 2	0	−4.6124	.0100
type 2	1	−1.0280	.6731
type 2	2	−1.2545	.5950
type 3	0	−7.3787	.0100
type 3	1	−2.7905	.2634
type 3	2	−1.7740	.6464
Phillips-Perron Test			
type 1	2	−1.7069	.4208
type 2	2	−24.9596	.0100
type 3	2	−34.4637	.0100
KPSS Test			
–	2	0.6486	.0182

TABLE A.5 Unit root tests for non-stationarity in external topics.

Non-stationarity type	*lag*	*statistic*	*p-value*
Augmented Dickey-Fuller Test			
type 1	0	−0.5879	.4583
type 1	1	0.1543	.6796
type 1	2	0.4255	.7577
type 2	0	−2.0858	.2997
type 2	1	−0.7813	.7582
type 2	2	−0.8569	.7321
type 3	0	−3.7039	.0426
type 3	1	−2.8528	.2408
type 3	2	−2.2273	.4682
Phillips-Perron Test			
type 1	2	−0.2539	.6237
type 2	2	−6.3929	.3547
type 3	2	−18.9020	.0381
KPSS Test			
–	2	0.7030	.0133

TABLE A.6 Regression tables: topic prevalence and citations.

	$log(c/y)$	$log[(c/y)+1]$	*hsin transf.*
log(Topic 3)	−0.209 (0.153)	−0.113** (0.055)	−0.146* (0.075)
Observations	675	777	777
R^2	0.522	0.690	0.663
Within R^2	0.984	0.983	0.982
log(Topic 11)	−0.069 (0.097)	−0.081* (0.048)	−0.107* (0.063)
Observations	675	777	777
R^2	0.464	0.644	0.618
Within R^2	0.982	0.981	0.980
log(Topic 19)	−0.190** (0.086)	−0.129* (0.070)	−0.170* (0.091)
Observations	675	777	777
R^2	0.494	0.650	0.625
Within R^2	0.983	0.981	0.980

Signif. Codes: ***: *0.01*, **: *0.05*, *: *0.1*. Standard errors clustered on year×journal level in parentheses. Further covariates: open access (Y/N), corresponding author, see eq. (1). If the corresponding author could not be identified using the Scopus database, the missing value has been replaced with a categorical value for the whole author team. Plot of the results in column 2 and 3 to be found in Figure 9. Separate regressions for each topic were conducted.

TABLE A.7 Unit root tests for non-stationarity in topics with negative citation correlation.

Non-stationarity type	*lag*	*statistic*	*p-value*
Augmented Dickey-Fuller Test			
type 1	0	−1.2871	.2108
type 1	1	−1.1637	.2545
type 1	2	−1.8260	.0677
type 2	0	−3.1899	.0353
type 2	1	−1.6696	.4497
type 2	2	−1.2875	.5836
type 3	0	−7.0824	.0100
type 3	1	−5.1028	.0100
type 3	2	−3.2884	.0933
Phillips-Perron Test			
type 1	2	−1.1818	.4677
type 2	2	−13.2054	.0416
type 3	2	−27.3217	.0100
KPSS Test			
–	2	0.7659	.0100

TABLE A.8 Unit root tests for non-stationarity of content similarity over time.

Non-stationarity type	*lag*	*statistic*	*p-value*
Augmented Dickey-Fuller Test			
type1	0	−0.9407	.3334
type1	1	−1.6127	.0982
type1	2	−2.0709	.0403
type2	0	−0.8835	.7229
type2	1	−0.9543	.6985
type2	2	−1.2101	.6103
type3	0	−3.7582	.0387
type3	1	−2.5902	.3363
type3	2	−2.2516	.4594
Phillips-Perron Test			
type 1	2	−0.5815	.5521
type 2	2	−1.7582	.7846
type 3	2	−16.8967	.0718
KPSS Test			
–	2	0.7185	.0119

TABLE A.9 Tests for cointegration.

Johansen test					
Coint.	**param.**	**LL**	**eigenvalue**	**trace stat.**	**5% crit. value**
0	6	−29.0492	–	19.7907	15.41
1¶	9	−19.6268	0.6103	0.9459	3.76
2	10	−19.1539	0.0462		
Engle-Granger test					
Regression Results First Step					
			Case Study Topics		
Competition Analysis			−2.920***		(0.653)
Constant			−6.139***		(0.811)
N			22		
$F(1, 20)$			20.02		
R^2			0.500		
adj. R^2			0.475		
Standard errors in parentheses					
$^*p < .05$, $^{**}p < .01$, $^{***}p < .001$					
Second Step: Dickey-Fuller test on residuals					
Non-stationarity type	***lag***	***statistic***	***adjusted DF statistic***		
			1%	**5%**	**10%**
type 1	1	−5.277	−4.29	−3.74	−3.45

Johansen Test: Constant trend; Lag length= 2, $N = 20$, time: 2000-2021. Engle-Granger test: $N = 22$, adjusted DF-statistic critical values based on (Davidson and MacKinnon, 1993, p. 722) since the standard DF-critical values cannot be used as the OLS estimation in the first step distorts the variance. ¶ marks the cointegration level at which the critical value exceeds the trace statistic.

TABLE A.10 Tests for cointegration for the two variable specification with three lags.

Johansen Test					
Coint.	**param.**	**LL**	**eigenvalue**	**trace stat.**	**5% crit. value**
0¶	10	−21.9278	–	14.6659	15.41
1	13	−15.5038	0.49146	1.8179	3.76
2	14	−14.5949	0.09124		

Johansen Test: Constant trend; Lag length= 3, $N = 19$, time: 2000-2021. ¶ marks the cointegration level at which the critical value exceeds the trace statistic.

TABLE A.11 Selection-order criteria VAR.

lag	**LL**	**LR**	**df**	***p***	**FPE**	**AIC**	**HQIC**	**SBIC**
0	−32.2897				0.1058	3.42897	3.4484	3.52854
1	−27.0995	10.38	4	.034	0.094303	3.30995	3.36826	3.60867
2	−19.1539	15.891¶	4	.003	0.064653¶	2.91539¶	3.01258¶	3.41325¶
0	−29.1206				0.0907	3.27586	3.29268	3.3753
1	−22.9299	12.381	4	.015	0.0724	3.0453	3.0957	3.3435
2	−15.5564	14.747¶	4	.005	0.0518¶	2.6902¶	2.7743¶	3.1872¶
3	−14.5949	1.9231	4	.750	0.0750	3.0100	3.1278	3.7059

¶ marks the best outcome in each column. For two lags: $N = 20$, for three lags: $N = 19$.

TABLE A.12 Test for normally distributed residuals of the VAR model (3 lags).

Jarque-Bera test				
Equation		χ^2	df	Prob $>\chi^2$
Case studies		0.198	2	.9058
Competition Analysis		0.361	2	.8346
All		0.559	4	.9675
Skewness test				
Equation	Skewness	χ^2	df	Prob $>\chi^2$
Case studies	−0.016	0.001	1	.9773
Competition Analysis	−0.311	0.306	1	.5799
All		0.307	2	.8576
Kurtosis test				
Equation	Kurtosis	χ^2	df	Prob $>\chi^2$
Case studies	2.501	0.197	1	.6571
Competition Analysis	2.736	0.055	1	.8144
All		0.252	2	.8815

TABLE A.13 Lagrange-multiplier test for autocorrelation of residuals in the VAR model.

lag	χ^2	**df**	**Prob** $>\chi^2$
1	8.6979	4	.069
2	6.8167	4	.146

H0: no autocorrelation at lag order *i*.

TABLE A.14 Eigenvalues for test of the Eigenvalue stability condition.

Eigenvalue		
real part	**imaginary part**	**Modulus**
0.8198		0.8198
0.0465	+ 0.6665*i*	0.6681
0.0465	− 0.6665*i*	0.6681
−0.6252	+ 0.1039*i*	0.6338
−0.6252	− 0.1039*i*	0.6338
0.2160		0.2160

All eigenvalues have a modulus <1. The VAR with 3 lags fulfills the stability condition. See Figure A.7 below for the unit circle.

TABLE A.15 Forecast error variance decomposition – table.

Step	FEVD	Std. Err.	95% CI	
0	0	0	0	0
1	0	0	0	0
2	0.3876	0.0556	0.7195	0.1694
3	0.5149	0.1789	0.8509	0.1714
4	0.5025	0.1545	0.8504	0.1775
5	0.4979	0.1536	0.8423	0.1757
6	0.4999	0.1520	0.8479	0.1775
7	0.5133	0.1604	0.8663	0.1801
8	0.5121	0.1551	0.8690	0.1821

Impulse variable: competition analysis, response variable: External topics. FEVD Computation for Figure 12.

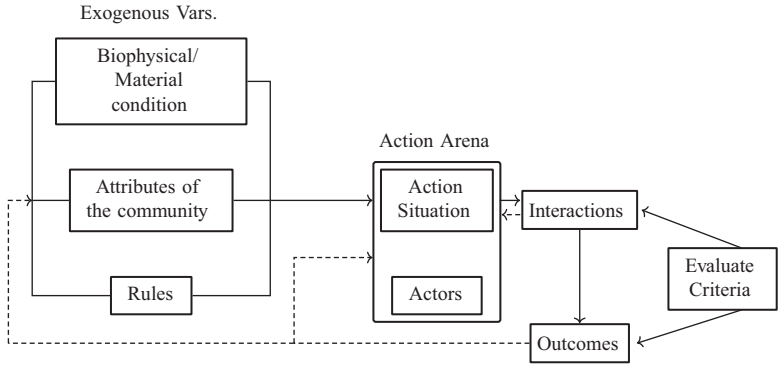

FIGURE A.1 The main structure of the IAD framework. Taken from (Ostrom, 2005, p. 15).

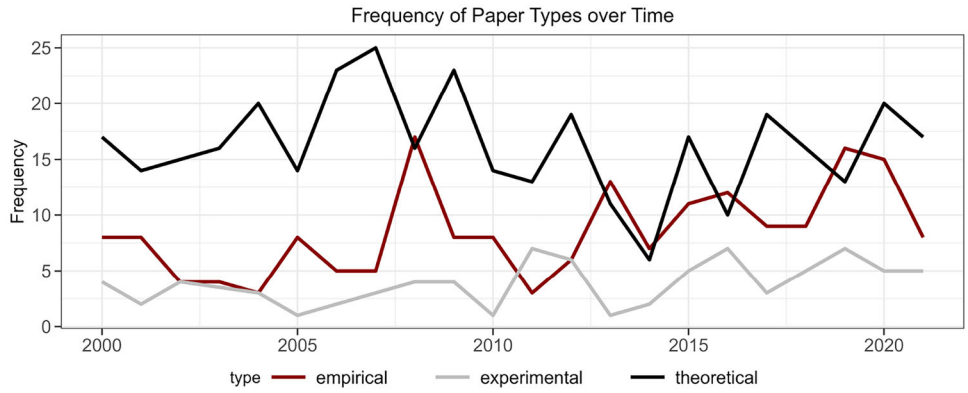

FIGURE A.2 Frequency of paper types over time. A less granular alternative to the use of latent topics is clustering papers solely by their methodological focus. This is less detailed as latent topics can also capture the methodology used and the content of a publication – a level of granularity a classification by method cannot provide. I distinguish between the categories theoretical, empirical, experimental, and policy, which I manually assign, that is, I handcode based on the abstract and, if necessary, the introduction whether a paper can be considered mainly theoretical, empirical, or experimental in its methodology. If a paper applies two categories, for example, by developing a theoretical model and testing it empirically, I review the particular publication and weigh the contributions to come to a conclusion. I do not cluster the antitrust publications in this step. I also omit policy papers as they only account for a very small fraction of the publications. This leads to the frequencies over time, as shown here.

One can see that there is a slight decline in theoretical papers in the years around 2010. However, reaching a minimum in 2014, the number of mainly theoretical contributions has begun to increase again. At the same time, one can observe a minor growth in empirical papers and a rather stable evolution of experimental ones. Even though there might be some changes over time, I mostly observe fluctuations, and changes in the topics' prevalence may have other reasons than shifts in the methodology applied. [Colour figure can be viewed at wileyonlinelibrary.com]

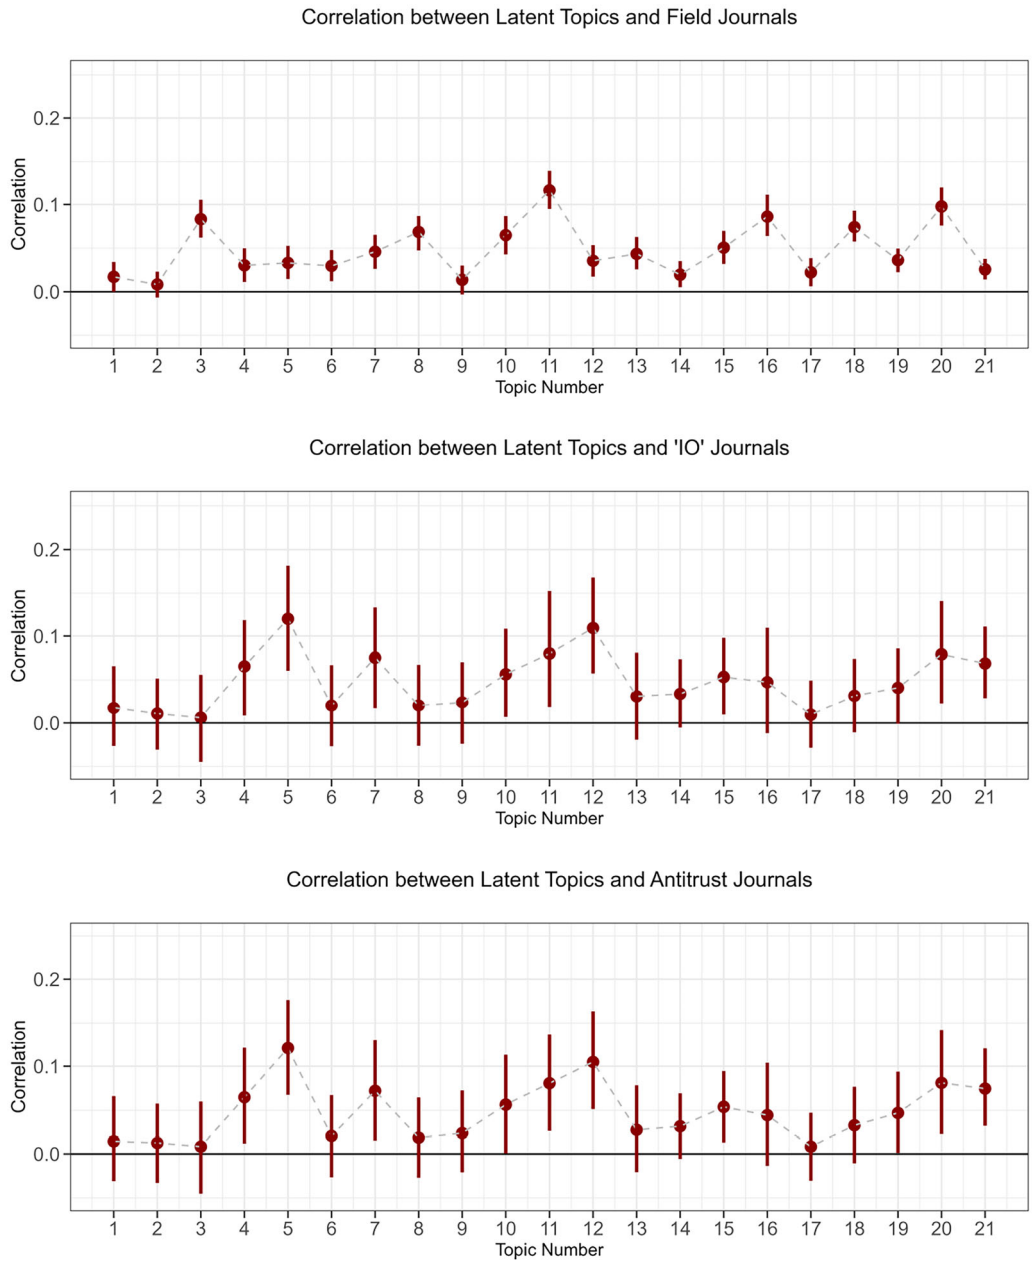

FIGURE A.3 Topic correlation with journal types. [Colour figure can be viewed at wileyonlinelibrary.com]

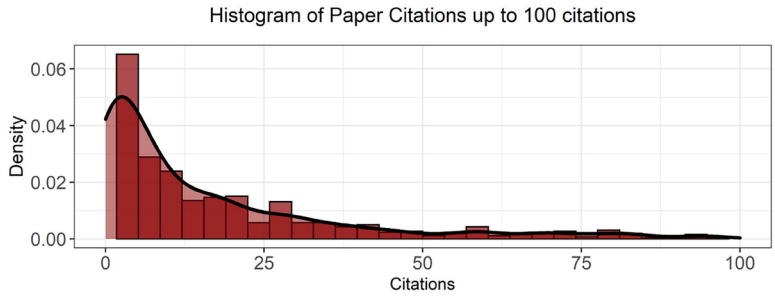

FIGURE A.4 Histogram of citations ≤ 100. [Colour figure can be viewed at wileyonlinelibrary.com]

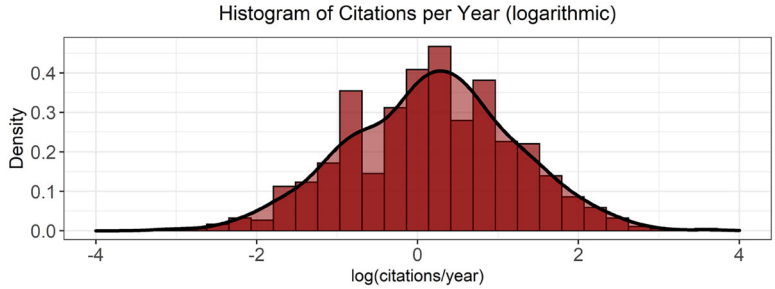

FIGURE A.5 Histogram of logarithmic citations per year. [Colour figure can be viewed at wileyonlinelibrary.com]

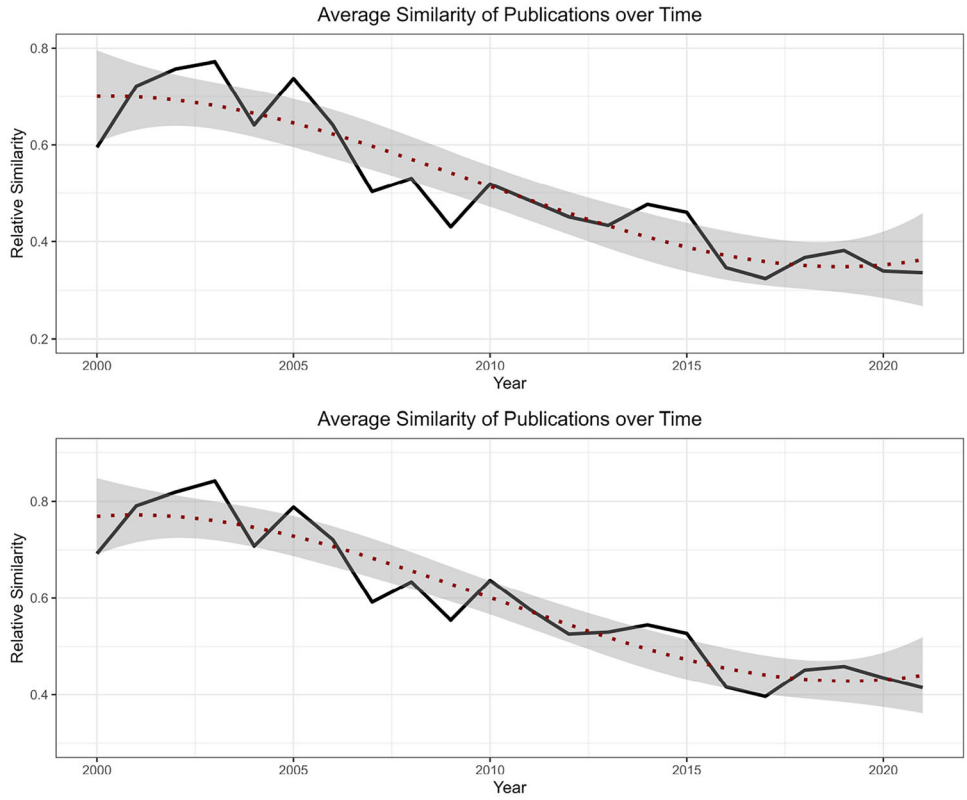

FIGURE A.6 Publication Similarity over time using different neural network specifications. Upper plot: Paragraph Vector dimension: 100. Lower plot: Vector dimension: 500. Method: paragraph vector with distributed memory (PV-DM), iterations: 20. [Colour figure can be viewed at wileyonlinelibrary.com]

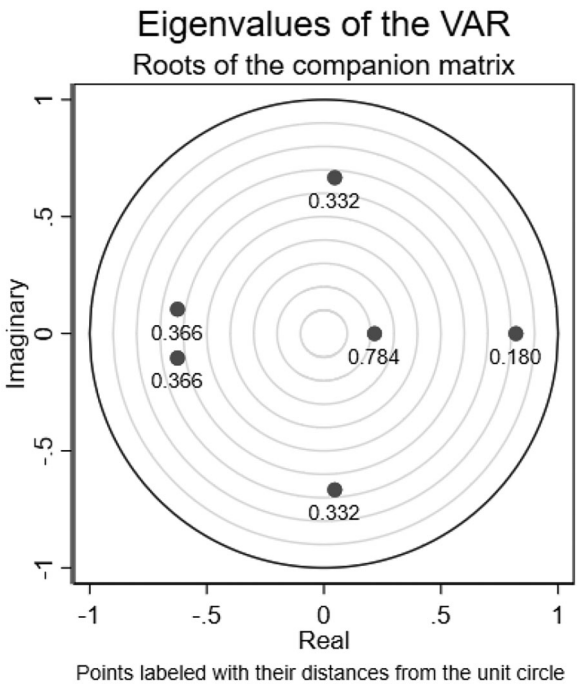

FIGURE A.7 Unit Circle of the VAR model.

APPENDIX B: SELECTION OF LATENT TOPICS

Figure B.1 presents the two selection measures exclusivity (EX) and semantic coherence (SC) for

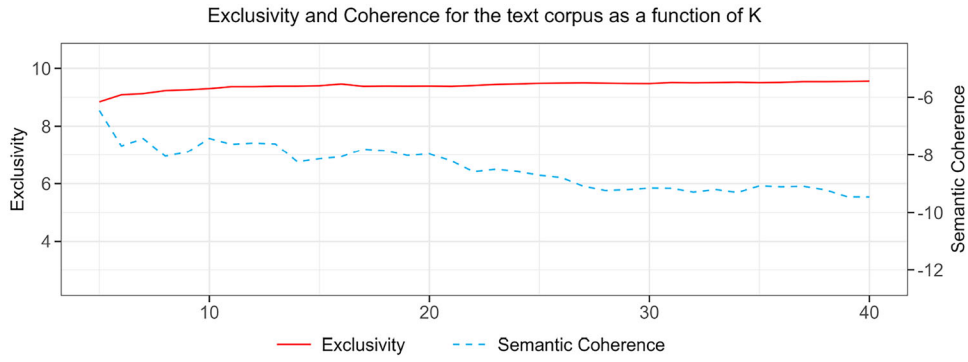

FIGURE B.1 Measures to identify the optimal K for the text corpus. $K \in [5, 40]$, Semantic Coherence adjusted: $\hat{SC} = \frac{SC}{10} + 16$ to be comparable with exclusivity. [Colour figure can be viewed at wileyonlinelibrary.com]

the text corpus. One can easily see a generally negative slope for semantic coherence and a generally positive slope for exclusivity as suggested in theory. The choice of the optimal K^* is subject to individual reasoning as mentioned beforehand. Still, I propose some attempt to formalize the decision based on EX and SC as follows:

$$SER_k = \frac{SC_k}{EX_k} \tag{B.1}$$

$$\Delta SER_k = SER_k - SER_{k-1} \tag{B.2}$$

$$wSER_k = \frac{\Delta SER_k}{SER_k} \tag{B.3}$$

Equation (B.1) describes the ratio of semantic coherence and exclusivity for each $K = k$. In a second step, I use the change in the SER to quantify the effect of a marginal (discrete) change in k to $k + 1$ as shown in eq. (B.2). This approach captures changes in both measures at the same time. As coherence and exclusivity are not stationary, eq. (B.3) weights the change in SER by the absolute value of the SER. Due to the nature of the data $EX > 0$ and $SC < 0$ holds, such that $SER_k < 0$ always holds. However, the change in ΔSER_k depends on how both measures jointly change. A sole decrease in exclusivity would imply $\Delta SER_k > 0$ as well as an increase in semantic coherence. Figure B.2 plots the measures shown in eqs. (B.2) and (B.3).

Eventually, I limit the number of possible topics to $K \in [5, 40]$ as due to the comparatively narrow scope of the paper selection I consider a sufficiently larger amount of separate topics captured in paper abstracts as unreasonable. On the other hand, it is unreasonable that only three or four underlying topics in collusion research exist. The absolute minimum of topics would be two as one topic would essentially cover everything.

Stepping back, one should bear in mind that the choice of K is up to the researcher and existing literature applying topic modeling algorithms to the best of my knowledge applies less sophisticated decision rules. Hence, the proposed decision rule attempts to contribute a new metric to find K^* but is still not independent from my personal understanding.

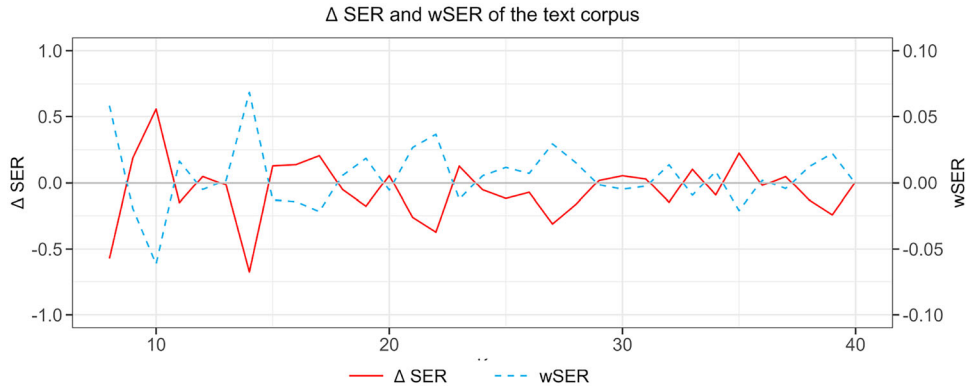

FIGURE B.2 ΔSER_k and $wSER_k$ for the text corpora. ΔSER_k (red solid line) and $wSER_k$ (blue dashed line) for the Text Corpus. $K \in [5, 40]$. [Colour figure can be viewed at wileyonlinelibrary.com]

Looking at Figure B.2, one can see that $wSER_k < 0$ often occurs, which corresponds to stronger fluctuations in semantic coherence across K values (see Figure B.1). Hence, semantic coherence increases occasionally with an increasing number of K. As one can see in Figure B.2, especially for $K < 30$ exists a lot of fluctuation. $K = 21$ appears as a reasonable amount of topics for the temporal range of 22 years and the chosen focus on collusive practices. However, close values such as 20 or 22 would be conceivable as well.

TABLE B.1 Latent topics in joint 21st century antitrust and collusion text corpus: Alternative choice of K.

Topic 1:	
Highest Prob:	competit, articl, author, enforc, law, agreement, polici
FREX:	crimin, articl, assess, remedi, programm, jurisdict, sanction
Topic 2:	
Highest Prob:	damag, case, law, court, competit, decis, european
FREX:	settlement, damag, court, platform, infring, suprem, appeal
Topic 3:	
Highest Prob:	market, firm, competit, industri, effect, model, product
FREX:	output, market, oligopoli, duopoli, share, joint, cournot
Topic 4:	
Highest Prob:	auction, bid, bidder, use, procur, test, two
FREX:	auction, bid, bidder, first-pric, ring, procur, revenu
Topic 5:	
Highest Prob:	equilibrium, inform, game, communic, firm, repeat, privat
FREX:	signal, equilibrium, communic, game, payoff, repeat, perfect
Topic 6:	
Highest Prob:	antitrust, competit, econom, industri, trade, polici, group
FREX:	athlet, ncaa, amateur, colleg, sport, centuri, ship
Topic 7:	
Highest Prob:	agent, incent, contract, princip, organ, optim, effort
FREX:	agent, princip, deleg, corrupt, effort, team, audit
Topic 8:	
Highest Prob:	price, firm, cost, consum, demand, model, market
FREX:	price, capac, demand, guarante, consum, war, airlin
Topic 9:	
Highest Prob:	lenienc, effect, firm, polici, fine, enforc, increas
FREX:	lenienc, program, fine, deterr, deter, complianc, prosecut
Topic 10:	
Highest Prob:	vertic, merger, retail, firm, supplier, buyer, market
FREX:	upstream, retail, vertic, supplier, merger, downstream, integr

Alternative value of $K = 10$. Text corpus contains 777 documents, 2590 terms and 37746 tokens. EM-Algorithm iterations: ≤ 75. Spectral initialization.

TABLE B.2 Correlation table of the topic model with $K = 21$ topics.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
1	1.00	0.02	−0.08	−0.07	−0.08	0.00	−0.08	−0.06	−0.02	−0.06	−0.11	−0.07	−0.05	−0.04	−0.06	0.00	−0.03	−0.06	−0.05	−0.10	−0.01
2	0.02	1.00	−0.09	−0.06	−0.08	−0.03	−0.07	−0.05	−0.04	−0.06	−0.12	−0.06	−0.04	−0.05	−0.04	−0.05	−0.02	−0.04	−0.04	−0.09	−0.05
3	−0.08	−0.09	1.00	−0.06	−0.08	−0.06	−0.08	−0.04	−0.07	−0.07	−0.04	−0.02	−0.04	−0.04	−0.05	−0.10	−0.04	−0.07	−0.04	−0.05	0.00
4	−0.07	−0.06	−0.06	1.00	0.03	−0.06	−0.05	−0.04	−0.05	−0.06	−0.09	−0.05	−0.06	−0.02	−0.06	−0.06	−0.04	−0.05	−0.02	−0.08	−0.04
5	−0.08	−0.08	−0.08	0.03	1.00	−0.08	0.01	−0.06	−0.07	−0.09	0.00	−0.04	−0.07	−0.06	−0.07	−0.08	−0.07	−0.06	−0.06	−0.04	−0.03
6	0.00	−0.03	−0.06	−0.06	−0.08	1.00	−0.06	−0.06	−0.03	−0.06	−0.10	−0.05	−0.05	−0.04	−0.03	−0.05	−0.02	−0.05	−0.04	−0.08	−0.03
7	−0.08	−0.07	−0.08	−0.05	0.01	−0.06	1.00	−0.07	−0.06	−0.07	−0.10	0.00	−0.06	−0.05	−0.07	−0.05	−0.06	−0.05	−0.03	−0.06	−0.05
8	−0.06	−0.05	−0.04	−0.04	−0.06	−0.06	−0.07	1.00	−0.04	−0.04	−0.04	−0.05	−0.04	−0.04	−0.06	−0.07	−0.03	−0.04	−0.05	−0.05	−0.05
9	−0.02	−0.04	−0.07	−0.05	−0.07	−0.03	−0.06	−0.04	1.00	−0.05	−0.08	−0.05	−0.05	−0.04	−0.05	−0.04	−0.03	−0.05	−0.04	−0.06	−0.02
10	−0.06	−0.06	−0.07	−0.06	−0.09	−0.06	−0.07	−0.04	−0.05	1.00	−0.06	−0.06	−0.02	−0.03	−0.03	−0.06	−0.03	−0.06	−0.05	−0.04	−0.05
11	−0.11	−0.12	−0.04	−0.09	0.00	−0.10	−0.10	−0.04	−0.08	−0.06	1.00	−0.03	−0.03	−0.06	0.03	−0.13	−0.09	−0.03	−0.07	0.00	0.03
12	−0.07	−0.06	−0.02	−0.05	−0.04	−0.05	0.00	−0.05	−0.05	−0.06	−0.03	1.00	−0.05	−0.05	−0.04	−0.08	−0.05	−0.03	−0.05	−0.05	0.00
13	−0.05	−0.04	−0.04	−0.06	−0.07	−0.05	−0.06	−0.04	−0.05	−0.02	−0.03	−0.05	1.00	−0.01	−0.04	−0.08	−0.03	−0.04	−0.04	−0.07	−0.05
14	−0.04	−0.05	−0.04	−0.02	−0.06	−0.04	−0.05	−0.04	−0.04	−0.03	−0.06	−0.05	−0.01	1.00	−0.04	−0.05	−0.04	−0.05	−0.02	−0.05	−0.02
15	−0.06	−0.04	−0.05	−0.06	−0.07	−0.03	−0.07	−0.06	−0.05	−0.03	0.03	−0.04	−0.04	−0.04	1.00	−0.04	−0.05	−0.05	−0.03	−0.09	0.04
16	0.00	−0.05	−0.10	−0.06	−0.08	−0.05	−0.05	−0.07	−0.04	−0.06	−0.13	−0.08	−0.08	−0.05	−0.04	1.00	−0.05	−0.05	−0.06	−0.08	−0.06
17	−0.03	−0.02	−0.04	−0.04	−0.07	−0.02	−0.06	−0.03	−0.03	−0.03	−0.09	−0.05	−0.03	−0.04	−0.05	−0.05	1.00	−0.04	−0.04	−0.06	−0.02
18	−0.06	−0.04	−0.07	−0.05	−0.06	−0.05	−0.05	−0.04	−0.05	−0.06	−0.03	−0.03	−0.04	−0.05	−0.05	−0.05	−0.04	1.00	−0.03	−0.07	−0.04
19	−0.05	−0.04	−0.04	−0.02	−0.06	−0.04	−0.03	−0.05	−0.04	−0.05	−0.07	−0.05	−0.04	−0.02	−0.03	−0.06	−0.04	−0.03	1.00	−0.05	0.00
20	−0.10	−0.09	−0.05	−0.08	−0.04	−0.08	−0.06	−0.05	−0.06	−0.04	0.00	−0.05	−0.07	−0.05	−0.09	−0.08	−0.06	−0.07	−0.05	1.00	−0.07
21	−0.01	−0.05	0.00	−0.04	−0.03	−0.03	−0.05	−0.05	−0.02	−0.05	0.03	0.00	−0.05	−0.02	0.04	−0.06	−0.02	−0.04	0.00	−0.07	1.00