

## RESEARCH ARTICLE

## Dropping diversity of products of large US firms: Models and measures

Ananthan Nambiar<sup>1,2</sup>, Tobias Rubel<sup>3,4</sup>, James McCaull<sup>5</sup>, Jon deVries<sup>4</sup>, Mark Bedau<sup>3\*</sup>

**1** Department of Bioengineering, University of Illinois at Urbana-Champaign, Urbana, Illinois, United States of America, **2** Carl R. Woese Institute for Genomic Biology, University of Illinois at Urbana-Champaign, Urbana, Illinois, United States of America, **3** Department of Philosophy, Reed College, Portland, Oregon, United States of America, **4** Department of Biology, Reed College, Portland, Oregon, United States of America, **5** Department of Computer Science, Reed College, Portland, Oregon, United States of America

\* These authors contributed equally to this work.

\* [mab@reed.edu](mailto:mab@reed.edu)

## OPEN ACCESS

**Citation:** Nambiar A, Rubel T, McCaull J, deVries J, Bedau M (2022) Dropping diversity of products of large US firms: Models and measures. PLoS ONE 17(3): e0264330. <https://doi.org/10.1371/journal.pone.0264330>

**Editor:** László Vasa, Szechenyi Istvan University: Szechenyi Istvan Egyetem, HUNGARY

**Received:** April 28, 2021

**Accepted:** February 8, 2022

**Published:** March 16, 2022

**Copyright:** © 2022 Nambiar et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** Data used to generate all figures is available at the following Github repository: <https://github.com/TobiasRubel/firm-diversity-data>.

**Funding:** This project was accomplished with the support of project "Predicting Technological Progress by Embedding Input-Output Networks in Semantic Technological Spaces" (PR0040), in response to IARPA-BAA-17-01 granted to Mark Bedau by IARPA (<https://www.iarpa.gov/>). The funders had no role in study design, data collection

## Abstract

It is widely assumed that in our lifetimes the products available in the global economy have become more diverse. This assumption is difficult to investigate directly, however, because it is difficult to collect the necessary data about every product in an economy each year. We solve this problem by mining publicly available textual descriptions of the products of every large US firm each year from 1997 to 2017. Although many aspects of economic productivity have been steadily rising during this period, our text-based measurements show that the diversity of the products of at least large US firms has steadily declined. This downward trend is visible using a variety of product diversity metrics, including some that depend on a measurement of the similarity of the products of every single pair of firms. The current state of the art in comprehensive and detailed firm-similarity measurements is a Boolean word vector model due to Hoberg and Phillips. We measure diversity using firm-similarities from this Boolean model and two more sophisticated variants, and we consistently observe a significant dropping trend in product diversity. These results make it possible to frame and start to test specific hypotheses for explaining the dropping product diversity trend.

## 1 Introduction

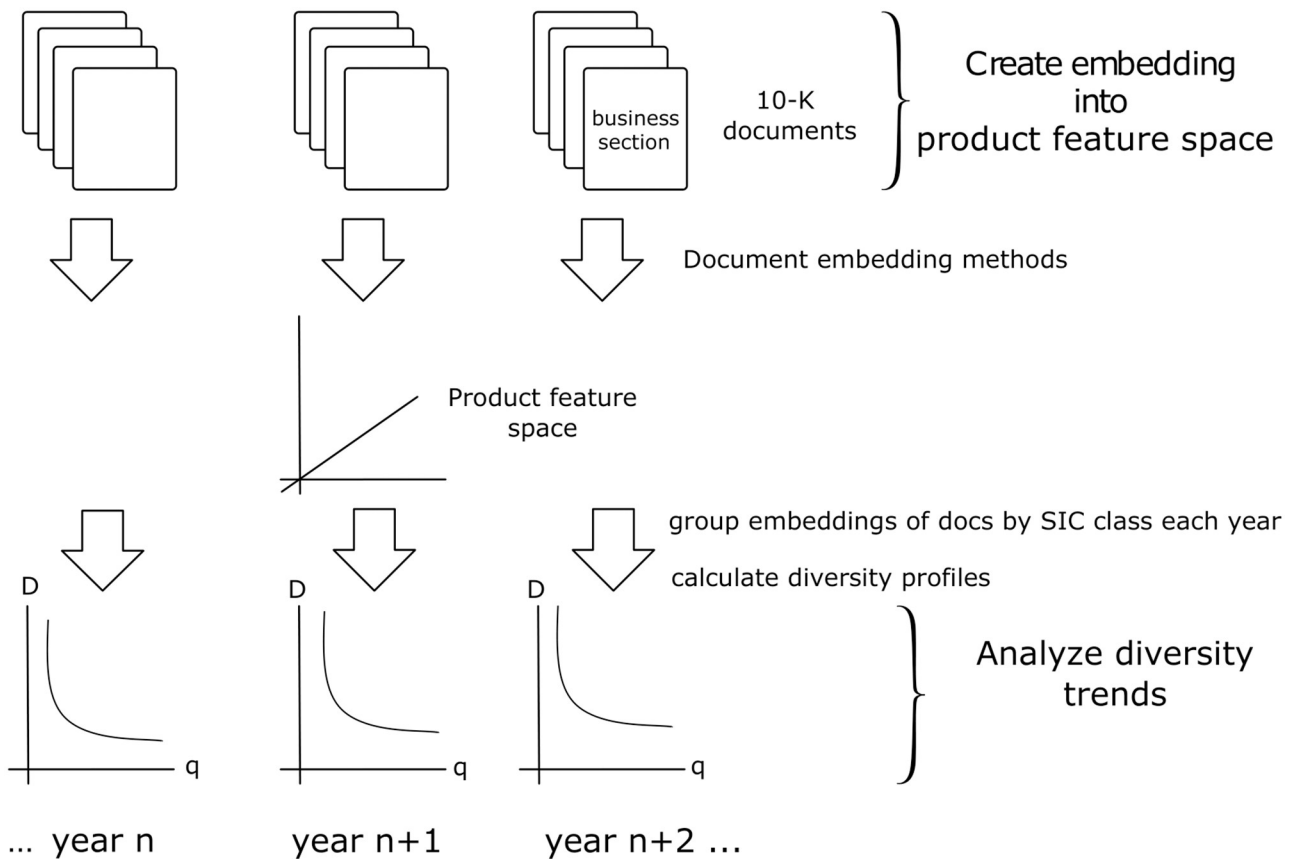
For decades economists have been using diversity to gauge the productivity and stability of regional economies, and this has motivated continuing efforts to craft better ways to measure diversity [1–5]. The economic diversity of geographic regions has been correlated with higher levels of gross domestic product, and economic diversification is often promoted as a route to economic stability, growth and development [6, 7]. This paper focuses more narrowly on the diversity of the *products* bought and sold in the economy overall. The diversification of products produced by important individual firms has been studied [8, 9], and so has the diversity of products in markets with many kinds of firms selling many kinds of products at fluctuating prices to many kinds of consumers [10–12]. Taking advantage of the existence of high

and analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** The authors have declared that no competing interests exist.

quality public textual data, this paper focuses on the products of large US firms over the past two decades.

Some discussions of product diversity are theoretical and focus on the mathematical consequences of simple economic scenarios, but our focus is empirical and data-driven, and relatively theory neutral and free of economic assumptions. We simply observe the changing product diversity of large US firms, evident in their annual product descriptions, and describe the trends we observe. In recent years, there have been similar efforts to draw ideas from quantitative biology, systems science and data mining to study the diversity of systems in social science [13, 14]. In economics, many papers design and apply standard indices of economic diversity and complexity (e.g. [15]), but atemporal data blinds us to how the indices have changed. The temporal data binning used here reveals how economic diversity and complexity have changed over the past generation and are trending today. By doing so, our results are precise and quantitative. In addition, our methods are easily reproducible. We first embed annual documents describing each firm’s products in a high-dimensional vector space, producing a model of the similarities among the products of large US firms. As shown in Fig 1, we then group the vectors by SIC class to obtain product-focused vector representations for industry classes. The diversity of those products is calculated from this classification for each year. We focus on three different document embeddings: a Boolean embedding modeled after the current industry standard in product-focused industry classification [16, 17], a slightly more sophisticated TF-IDF embedding, and a more complex Paragraph-Vector Distributed Memory (PV-DM) embedding. All of the models are first evaluated by



**Fig 1. The workflow for a text-based analysis of product diversity.**

<https://doi.org/10.1371/journal.pone.0264330.g001>

measurement of their Industry Specificity relative to the Standard Industrial Classification (SIC) and evaluation of the *a priori* plausibility of their firm clusters. Models that pass these tests are each used to measure the diversity of the products of large US firms over the past two decades. In order to identify diversity trends that are robust, we employ a suite of more or less complex ways to measure diversity, including a baseline measurement based merely on each firm's SIC classification. This enables us to identify diversity trends that are robust across a variety of models.

By doing so, we provide evidence during 1997–2017 of a falling trend in the diversity of products offered by large US companies. This evidence comes from a consensus of semantic-vector models trained on a corpus of 10-K documents from 1995–2019 that describe the products of those firms. This trend is further corroborated by the text-free model based just on SIC Codes. We conclude by evaluating a number of hypotheses for how to explain the trend of dropping diversity.

Our work is one of a growing number of text-based analyses of economic topics, such as banking, finance, accounting, mergers and acquisitions, or corporate innovation and fraud. Many use topic modeling methods akin to our methods [18–21] and apply them like we do to 10-K documents [18–20], while others mine other kinds of documents, such as IPO prospectuses [22–24] and analysts' reports and regulatory filings [23, 25–27].

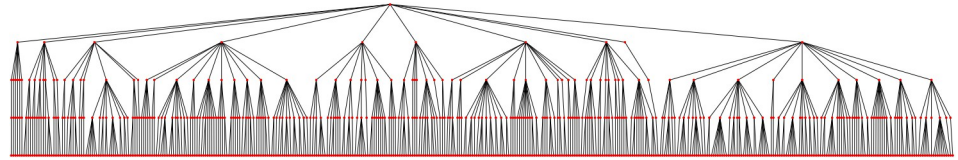
Our work also reflects the expanding diversity of applications of NLP and machine learning methods. Bergeaud and colleagues [28] used a similar methodology to classify patents by training models on patent documents, and their success motivates our application to economic documents of a more sophisticated methodology that was recently used to visualize and quantify the open-endedness of the evolution of technology [29]. Very similar computational methods were also recently used to infer models of the periodic table of elements from a training corpus of chemical documents [30]. These methods have even successfully predicted the biological function of a protein from its amino acid sequence, by training models on a huge corpus of amino acid sequences [31–33]—a vivid demonstration of the power and generality of our methods.

## 2 SIC model of firm similarity

Many measures of diversity operate on kinds or classes of things. Firms are regularly classified into two hierarchical classifications: The Standard Industrial Classification (SIC) [34] used by the Security and Exchange Commission (SEC) and the more recent North American Industry Classification System (NAICS) [35]. Both classifications were manually designed by experts and are updated by hand as industries evolve. In general, the NAICS classifies companies according to the processes by which they produce products, while the SIC classifies them according to the types of products they produce [35]. Given our present purpose of measuring diversity of products, this paper uses the SIC classification of firms when measuring the diversity of their products.

In a hierarchical classification tree like the 4-digit SIC classification scheme, individual firms  $i$  and  $j$  are leaves at the bottom of a 4-level branching tree structure. For example, the SIC hierarchical classification tree depicted in Fig 2 has 10 Divisions at the top level right below the tree's root, 83 Major Groups at the second level, 248 Industry Groups at the third level, 399 Codes at the fourth level; each individual firm is classified and under exactly one SIC Code. Each large US firm has a 4-digit SIC Code that specifies the firm's Division, Major Group, Industry Group, and Code.

Fig 3 shows the number of large US firms (in our training corpus) over the past twenty years. We can track changes in the number of firms in each SIC Division because firms in the

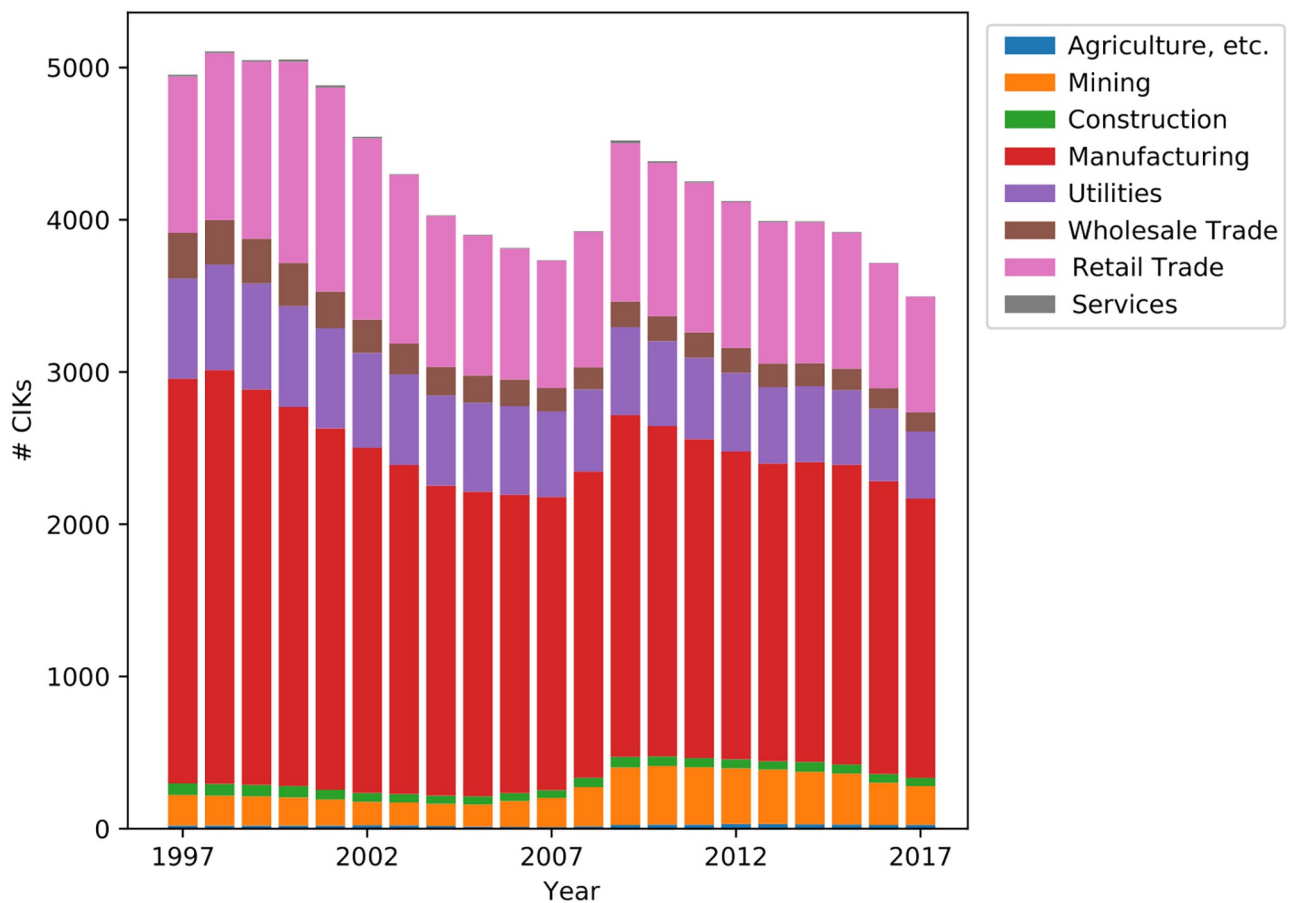


**Fig 2.** The SIC hierarchical classification tree of the documents in the training corpus, with Divisions (top level), Major Groups (second level), Industry Groups (third level), and Codes (bottom level).

<https://doi.org/10.1371/journal.pone.0264330.g002>

Figure are colored by their SIC Division. We can see that most SIC Divisions have shrunk somewhat so far in this century, although SIC Divisions typically retain roughly the same fraction of firms. The size of the Manufacturing Division dominates the pool, followed by Retail Trade and Utilities. The smallest Divisions are Services (almost invisible at the bottom of the bars), Construction, and Wholesale Trade. The only Division that shows significant growth in the past twenty years is Mining, which ended much larger than it started.

A simple gauge of the similarity of two firms is their *distance* from one another in the four-level SIC classification tree. We define the *distance* between firms  $i$  and  $j$  as the length of the shortest tree walk (sequence of adjacent nodes) between leaves  $i$  and  $j$ . The number of sub-classes in the SIC classification tree varies significantly across the different nodes in the tree. To



**Fig 3.** Annual count of 10-K documents in the training corpus; colors indicate SIC Divisions.

<https://doi.org/10.1371/journal.pone.0264330.g003>

create more distance between firms classified under especially heavily branching nodes, we define the length of a walk as the number of SIC codes that fall under its highest node.

The Standard Industrial Classification (SIC) tree has been carefully designed by human experts; it has passed the test of time and is widely used. We use it here to define a simple trustworthy metric of firm similarity against which to compare more sophisticated alternatives. This firm similarity metric based on proximity in the SIC classification tree is a crude representation of the similarities of actual firms. For example, the SIC tree proximity metric assigns a perfect similarity to every pair of firms in the same SIC Code, and it assigns identical similarities to all pairs of firms connected through the same highest node. This metric has a perfectly simple and predictable form, consisting of a number of rectangular fields with absolutely uniform similarity (Fig 4).

Embedding firms in semantic vector spaces provides a much more sensitive and product-centric measure of firm similarity. Each individual firm has a unique location in the vector space, which yields a fine-grained measure of the similarity of each pair of firms. The current industry standard in precise firm similarity matrices for large US firms is a simple Boolean word-vector embedding of documents [16, 17]. We construct and study an analogous Boolean word-vector model of product similarities, and we also construct and study two more sophisticated vector spaces. After confirming the plausibility of all of the models, we examine what they reveal about trends in the diversity of products of large US firms.

### 3 Semantic vector model-training corpus

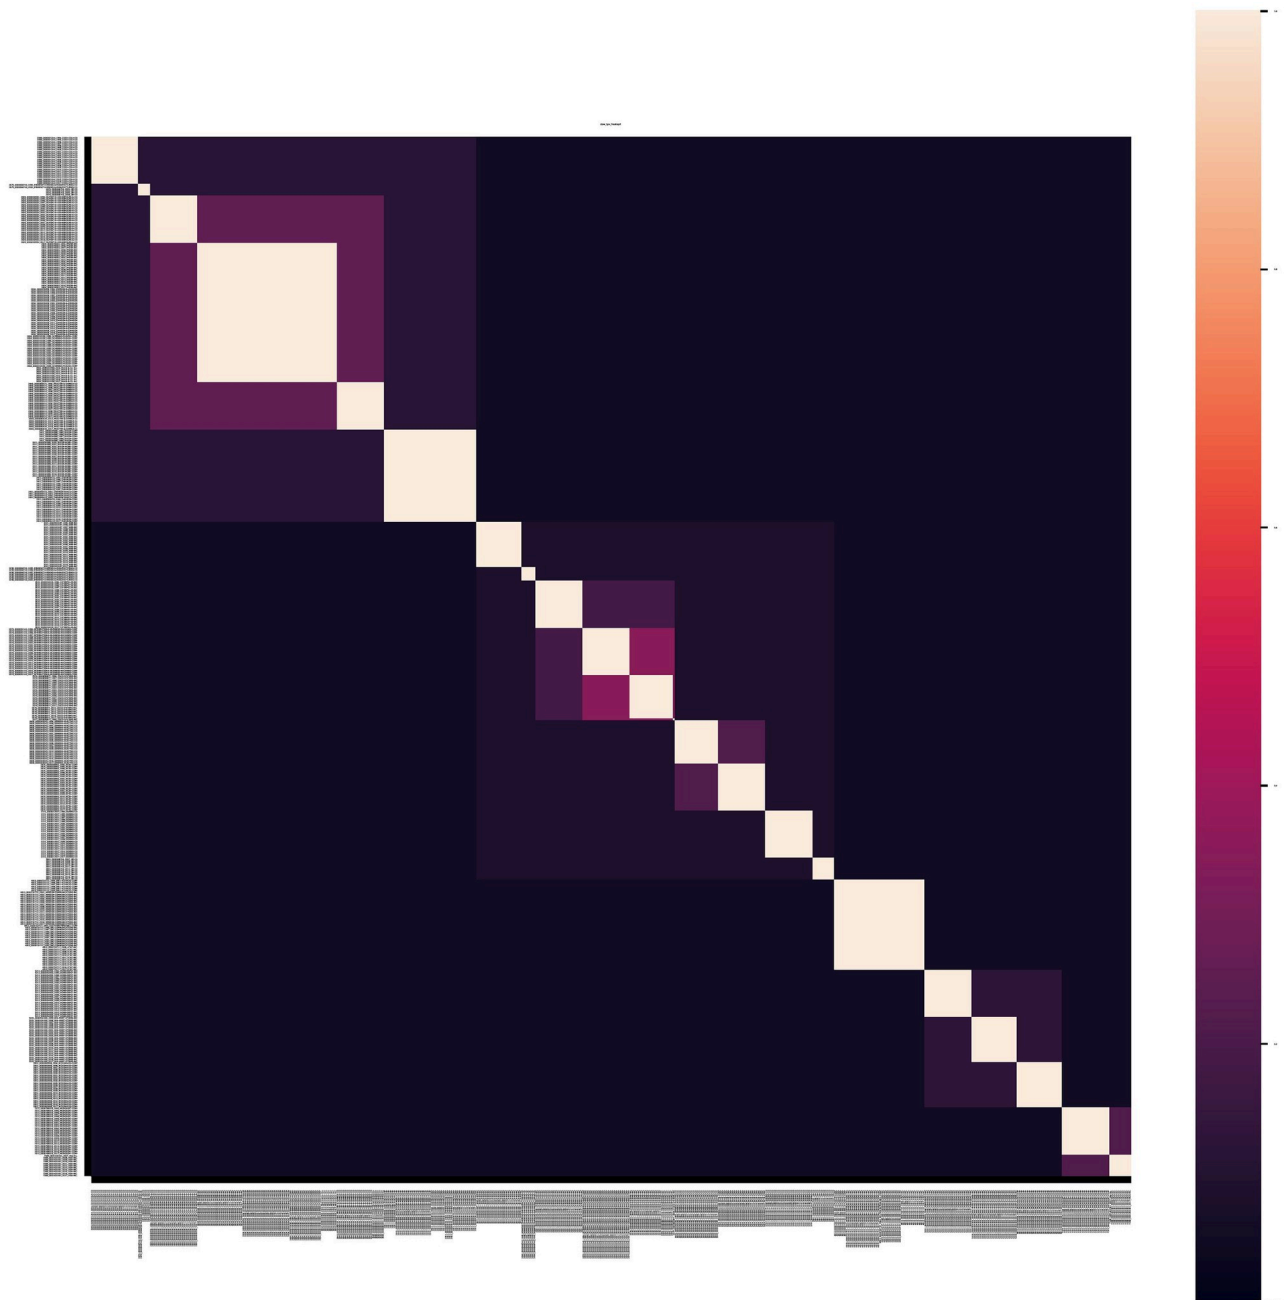
In order to build the product vector space we use the Form 10-K, a document filed with the SEC by any company with more than \$10 million in assets with ownership by 2000 or more individuals. The 10-K filing “provides a comprehensive overview of the company’s business and financial condition” [36]. Companies that file 10-K forms with the SEC are large US firms. Taken together the 10-K corpus is a complete, accurate, standardized, publicly available annual description of the products produced by every large US firm, and it was used to train the current industry standard in quantitative firm similarity measurement [16, 17].

We use the section of 10-K documents typically labeled “Part 1 Section 1: Business”. The Business section of a firm’s 10-K describes significant products the firm offers to their customers, what markets the firm operates in, and any subsidiaries it owns [16, 17]. If it exists, we exclude the part of the Business section typically labeled “Section 1A: Risk Factors,” leaving only details relevant to offered products.

We obtain 10-K, 10-K405, and 10-KSB documents from 1993 through 2018 from the Software Repository for Accounting and Finance (SRAF) [37]. The 10-K documents do not all have one standardized format, and their heterogeneity makes it a challenge to extract exactly their Business sections. SRAF stage-one parsing removes various markup from the documents and removes tables. Fig 3 shows the number of unique companies which file for each year in our dataset (broken down by SIC division).

After obtaining the data we extract the desired section by way of a series of regular expressions designed to catch common formats as well as a more flexible keyword based program. In total approximately 12% of documents cannot be parsed by either the regular expression or the keyword approaches. As Fig 5 illustrates, filing data for each company exists for only a subset of the years considered, but in general our programs are able to extract business sections from filings whenever the filings exist.

We evaluated the success of our extraction by manually checking both the extracted business sections to ensure that they were complete and did not contain extra text, as well as by reading through the unparseable documents to see if there was actually business section

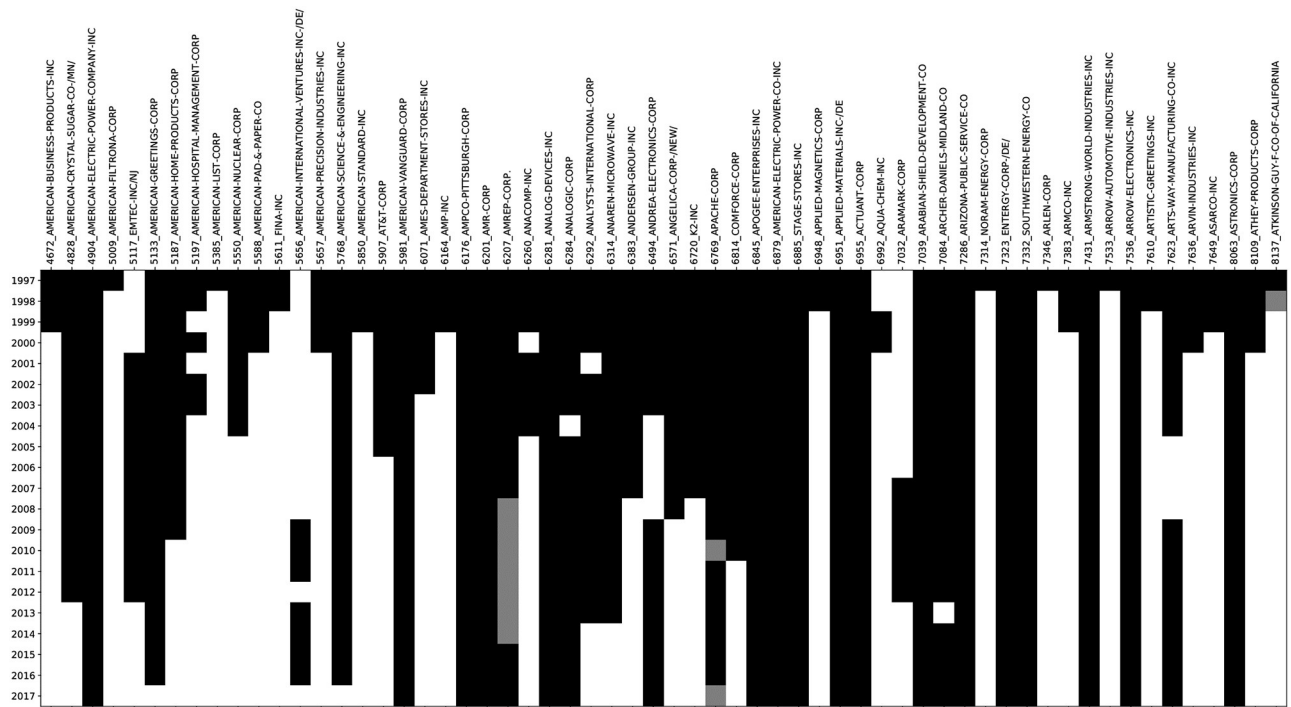


**Fig 4. Heatmap of the similarity matrix among twenty years of products products by the firms in the DJIA.** Rows (columns) of firms are ordered first by their 4-digit SIC Codes, and then by the CIK of each firm. The white squares along the diagonal show that pairs of documents with the same SIC Code are interpreted as having perfect similarity, and the fainter rectangles falling away from the white squares show that pairs of firms from coarser-grained SIC groups have a uniform intermediate degree of similarity.

<https://doi.org/10.1371/journal.pone.0264330.g004>

information lost by excluding those filings. Analysis of 50 randomly chosen extracted business sections revealed 49 of them to be correctly pulled from the corresponding 10-K forms. The errant filing was such that sections beyond the business section were included in the extracted text. Manual analysis of 100 randomly chosen unparseable filings found that 90 of them contained no business section at all, while the other 10 had either especially non-standard





**Fig 5. Map of annual documents for 30 firms in the model training corpus.** Documents are sorted into firm (columns) and filing year (row). Black cells indicate documents in the corpus, grey cells indicate documents in the SRAF dataset that were removed because they lacked the information we needed. White cells indicate years when there is no document for that company in the SRAF dataset.

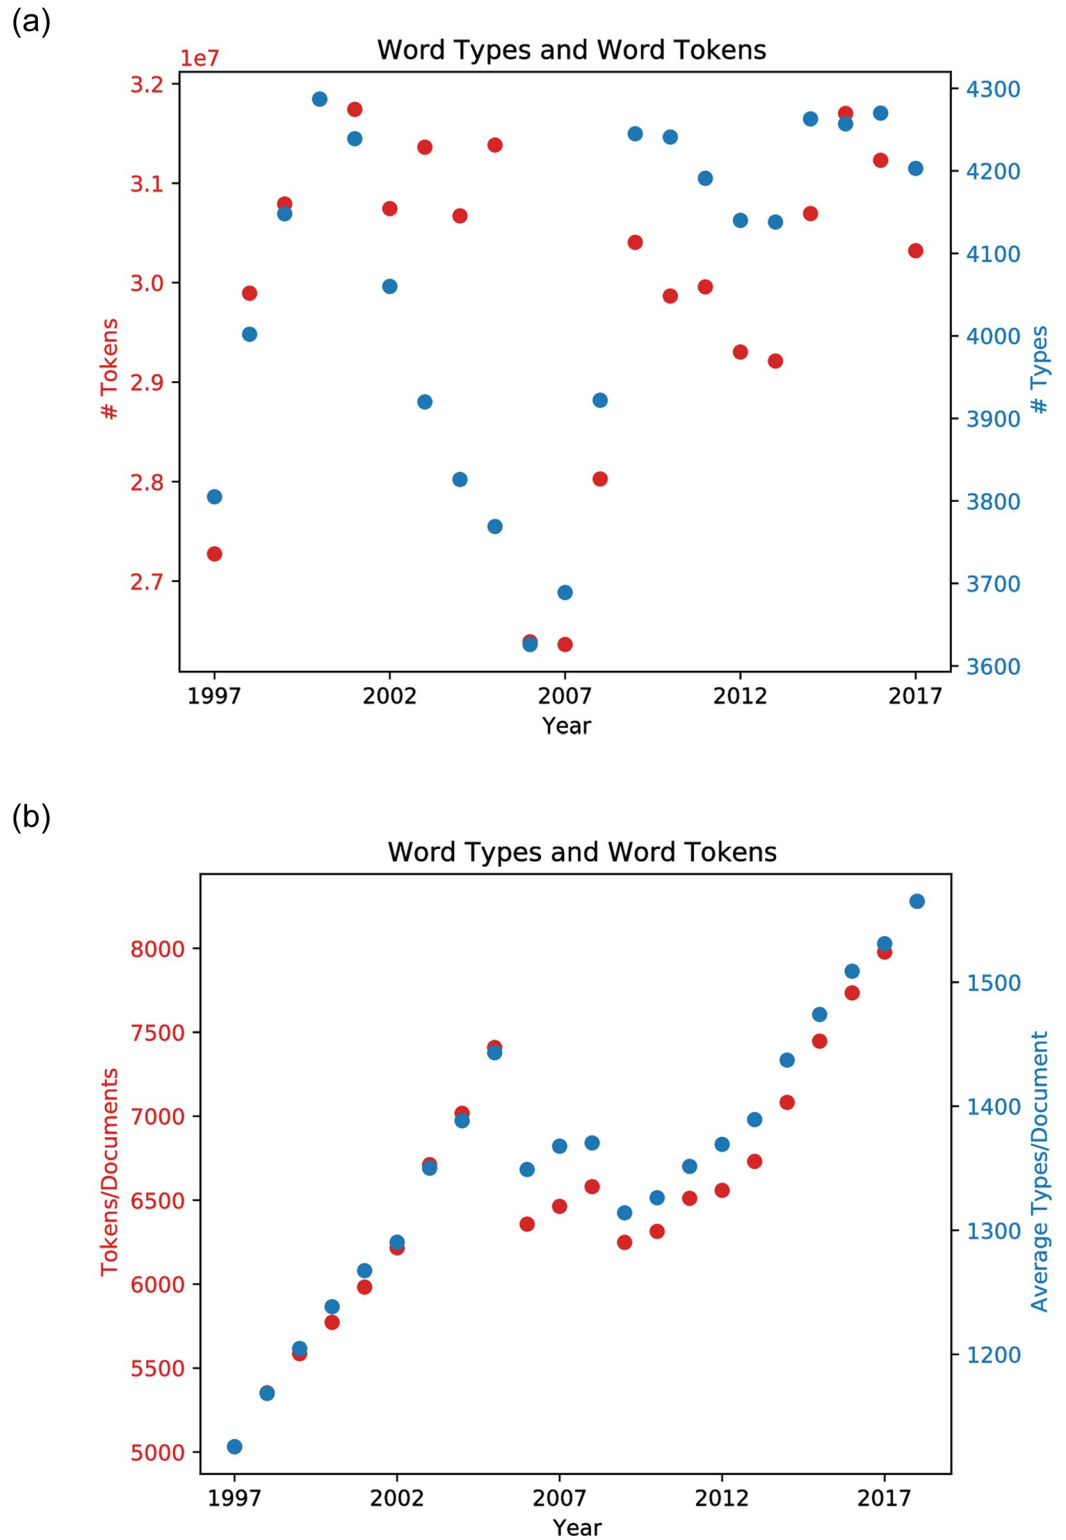
<https://doi.org/10.1371/journal.pone.0264330.g005>

formatting, extremely short business sections of less than 1000 characters, or combined their business and properties sections into a single section which made the relevant details of the section hard to distinguish from the irrelevant details. These analyses make us confident that we are building models on a dataset which is reasonably complete as well as textually relevant.

Once the appropriate sections are extracted they are preprocessed to only include nouns as suggested by [17]. In addition, we convert all alphabets to lower-case, remove extra white spaces, numeric values, stop words and words shorter than 3 characters long. In order to facilitate comparison with [17], we also remove from the training corpus any filings which do not have Compustat Global Company Keys, which lack at least a year of lagged Compustat data, or which are financial firms (SIC Codes 6000–6999), again following [17]. While the notion of a product can be extended to include some of the things that are “produced” by some financial firms, many large US financial firms do not offer the consumer products on which our analysis focuses. This last step reduces the number of individual documents in our training corpus from 179,717 to 107,500.

The number of different CIKs in the 10-K documents filed each year is plotted in Fig 3. This plot shows the size and Divisional composition of the documents used to train semantic vector models. The figure shows that across the years some SIC Divisions consistently contain at least an order of magnitude more 10-K forms than some other Divisions, with Manufacturing always dominating and three other Divisions (Agriculture etc., Construction, and Wholesale Trade) always lagging far behind. The figure also shows a largely consistent downward trend in the number of firms, with an exceptional uptick during 2008 and 2009.

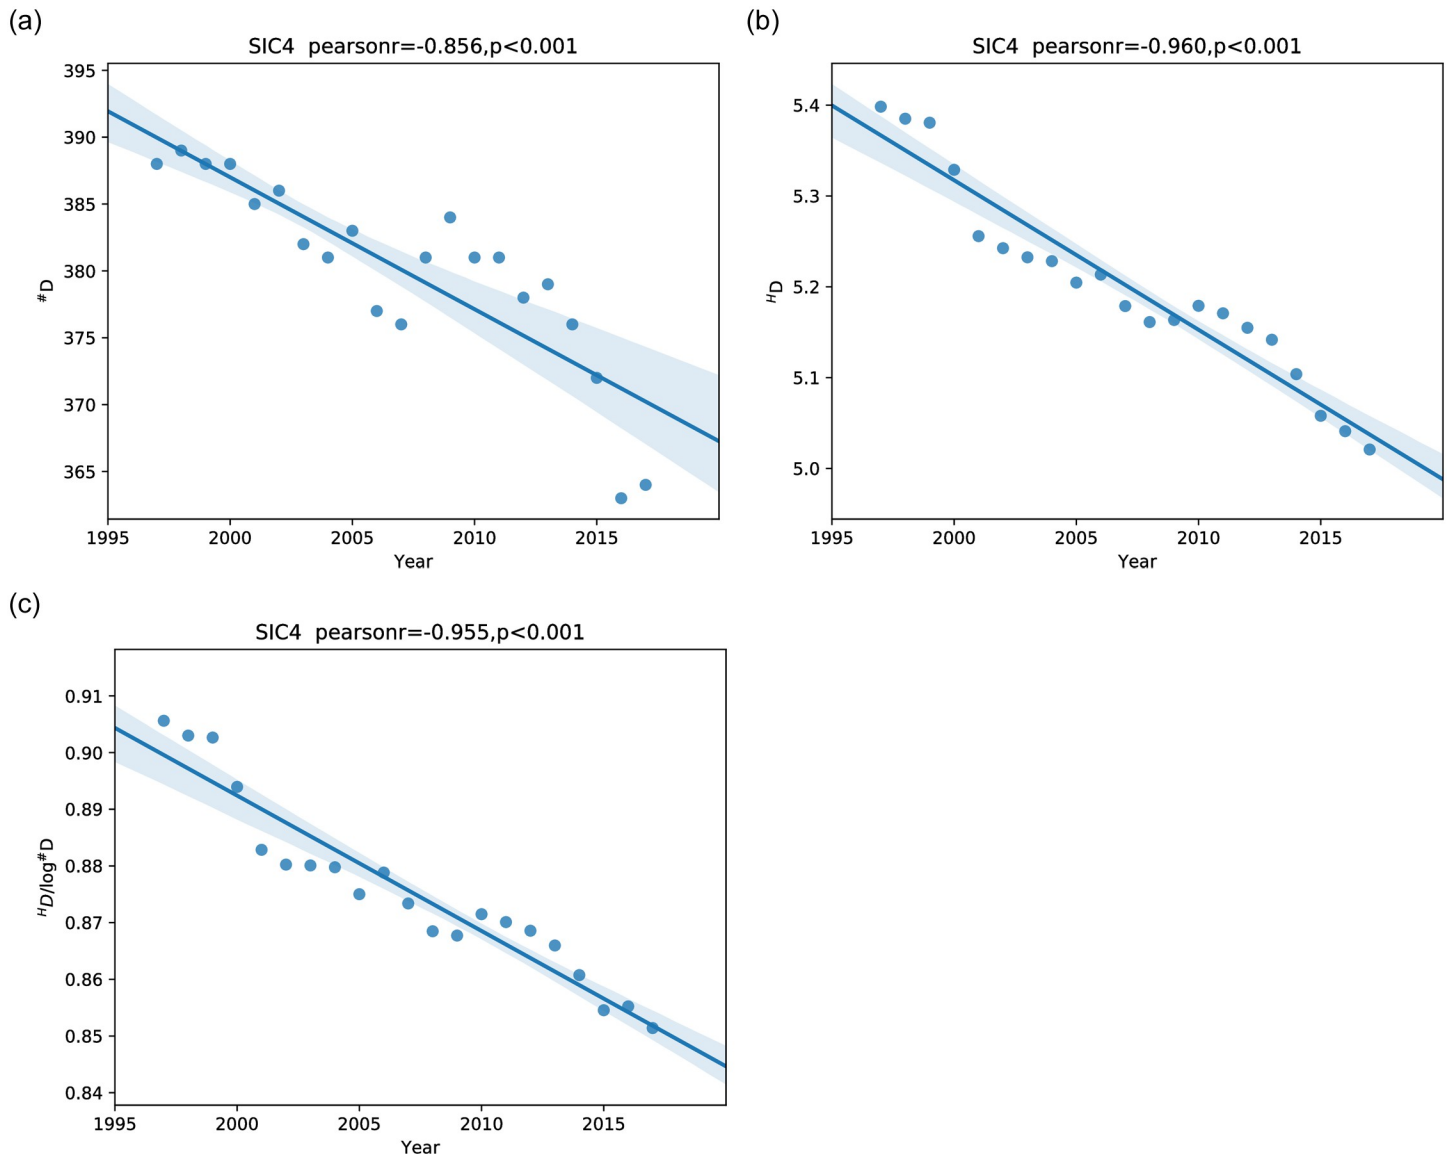
Text documents contain a number of *types* of words, and each type of word might have many different instances or *tokens* in a document. The left of Fig 6 shows the total number of



**Fig 6.** Left: The number of word tokens and word types in annual batches of 10-K documents in training corpus, during the first two decades of this century. Right: The mean number of word tokens and types in each document in each batch, over the same years.

<https://doi.org/10.1371/journal.pone.0264330.g006>





**Fig 7.** Left: The number of the 4-digit SIC Codes that are instantiated, during the first two decades of this century. Middle: Shannon entropy of the distribution of document counts per SIC Code during those years. Right: Normalized Shannon entropy of the same distribution during those years. Each plot includes a scatter plot of diversities and a linear regression fit with its 90% confidence interval.

<https://doi.org/10.1371/journal.pone.0264330.g007>

word tokens and types in the training corpus each year, and the right of the figure shows the average number of word tokens and types in each document each year. Aside from a drop during 2006 and 2007, the total number of word tokens and types are each fairly consistent over-age. By contrast, the *mean* number of word tokens and types per document rose significantly. (The mean also shows an analogous drop during 2006–2009.) The mean number of word tokens in a training document grew sixty percent during the first two decades of this century, and the number of word types grew more than forty percent.

Further analysis of the annual batches of 10-K documents in the training corpus shows that the number of instantiated SIC Codes has generally dropped during the first two decades of this century (Fig 7, left). At the start of the century an annual batch of 10-K documents covered firms from 392 SIC Codes, and twenty years later that number has dropped to 263 Codes.

Fig 7 (middle) shows twenty years of the Shannon entropy of the distribution of the number of instances of each SIC Code, revealing a general downward trend. The entropy of a distribution reflects both its width (number of bins) and evenness (similarity of counts across all bins). Since the number of bins (instantiated SIC classes) varies by more than 7% across the years we studied, it is also interesting to plot just the evenness of the distribution, which is shown by the normalized entropy:

$${}^H D_{adj} = \frac{{}^H D}{\log^{\#} D} \tag{1}$$

where  ${}^H D$  is the Shannon entropy and  $\#D$  is the number of bins (Fig 7 (right)). The distribution’s entropy and its width and evenness all display decreasing trends.

The number of instantiated SIC Codes is a simple measure of the diversity of the products produced by large US firms, as are the normalized or non-normalized Shannon entropy of the distribution. But both product diversity metrics are crude, because they ignore the different degrees of similarity between different SIC classes.

### 4 Embedding product descriptions with models

The documents in the training corpus are used to train a firm-similarity model that contains a vector representation of the products of each firm. Specifically, for every document  $p \in F$ , the embedding function is given by  $f_e: p \rightarrow v_p$  where  $v_p \in \mathbb{R}^d$ . All these vectors are normalized to have a length of 1. Here we compare embedding methods for bag-of-words models and neural network models.

#### 4.1 Bag-of-words embeddings

Bag-of-words models ignore the order of words in the training corpus and build vectors based just on the occurrence of the words. We study two different bag-of-words embeddings: Boolean and Term Frequency-Inverse Document Frequency (TF-IDF).

In the Boolean model, the vector for document  $p$ ,  $v_p$  is given by

$$v_p[i] = \begin{cases} 1, & \text{if } \Sigma[i] \in p \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

for every word  $\Sigma[i]$  in the dictionary. Following [17], a word is included in the dictionary only if it appears in less than 20 percent of the documents in the training corpus. Removing very common words is important but it is arbitrary to set a threshold at precisely 20 percent.

A more principled method is to replace the Boolean information about a word with the word’s term frequency/inverse document frequency (TF-IDF) statistic—a commonly used measure of the relevance of each word in a document from a large corpus. The TF-IDF model defines  $v_p$  with

$$v_p[i] = \begin{cases} \text{count}(\Sigma[i], p) * \log\left(\frac{|F|}{\text{docs}(\Sigma[i], F)}\right), & \\ \text{if } \Sigma[i] \in p & \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

where  $\text{count}(\Sigma[i], p)$  counts the number of times  $\Sigma[i]$  occurs in document  $p$ ,  $\text{docs}(\Sigma[i], F)$

counts the number of documents  $\Sigma[i]$  occurs in and  $|F|$  is the total number of documents in the corpus [38].

## 4.2 Neural embeddings

To obtain neural embeddings of firms in product space we use the Paragraph Vector Distributed Memory (PV-DM) model originally introduced in [39]. For a window size of 3, given words  $a_i$ ,  $a_{i+1}$  and  $a_{i+2}$  from document  $p$ , a neural network (NN) is trained to predict  $a_{i+3}$ . The PV-DM model then defines  $v_p$ , the vector representation of document  $p$ , as the  $d$ -dimensional hidden layer of the NN that is obtained from the document token. We omit details of the PV-DM architecture here since it is widely used and was followed here in all important respects. As suggested by the original paper [39], we set the number of dimensions of our PV-DM model at  $d = 300$ , and we trained the model for 20 epochs with a starting learning rate of 0.025 (which decays linearly between epochs) and a window size of 8. To train our PV-DM model, we take advantage of the Gensim library developed by [40].

## 5 Methods of analysis

Before we use our models to make more sophisticated measurements of the diversity of the products, we first establish the plausibility of the embeddings of firms in semantic vector spaces produced by the Boolean, TF-IDF, and PV-DM models. We gauge model plausibility in two ways: One is to measure how much similarity the embeddings attribute to the firms within the same industries, where the industries are identified by some trusted source. The other is to examine whether the micro-structure of the embeddings fit with human common sense judgments of the similarity of well-known firms.

### 5.1 Industry specificity

Existing classifications such as the SIC consider firms in the same industries to be relatively similar, and firms in distinct industries to be much less similar. The SIC is constructed by domain experts and is widely used by researchers and government offices, so it is safe to assume that each industry defined by a 4-digit SIC Code contains firms that are rather similar, much more similar than firms with different SIC Codes. So, one way to assess the plausibility of the vector embeddings of documents by individual firms is simply to check whether the average similarity of pairs of documents from firms in the same SIC Code is much higher than the average similarity of firms with different SIC Codes. The ratio of these two averages we term the *Industry Specificity (relative to the SIC)* of the similarity matrices produced by a given model. (See S1 Appendix in [S1 File](#) for precise definitions.)

### 5.2 Diversity

*Diversity* of products is often measured in economics simply as the number of different types of commodities (goods, products) available in a marketplace [10–12]. This approach is roughly analogous to the plot in of the number of different SIC Codes exemplified each year by large US firms (Fig 7 left). Sometimes the distribution of types of commodities in a market is weighted in some way, such as by total sales, and diversity is then measured by something like the Shannon entropy of the distribution of products [8, 9], an approach analogous to the Shannon entropy of the distribution of SIC Codes instances shown in Fig 7 (right). This entropy measure is quite simple, but is also rather crude, too crude, for example, to reflect the diversity of the firms within each SIC Code, or the “distance” between different SIC Codes within a given SIC Industry Group.

A more fine-grained approach is to measure the variance of the vectors in a product feature space by computing the number of dimensions needed to account for the bulk (here, 90%) of the variance of all of the individual firm vectors in each year. This measure has the virtue of being built out of the local details of the embedding of firms in a product space, and the results are relative to that product space. This measure is easily applied to documents that have been embedded in any product space of interest, and here we use the Boolean, TF-IDF, and PV-DM vector spaces.

An even more fine-grained measure of the diversity of the products produced by a set of firms comes from a generalized measure of diversity from theoretical ecology. Once a classification with  $s$  classes (as defined by four-digit SIC Codes) is obtained for a particular year  $y$ , a similarity matrix between different classes  $Z_y$  and a normalized abundance vector  $a = [a_{y,1}, \dots, a_{y,s}]$  is calculated. The diversity is then defined as

$${}^q D(a_y, Z_y) = \left( \sum_i a_y[i] \left( \sum_j Z_y[i, j] a_y[j] \right)^{q-1} \right)^{\frac{1}{1-q}} \quad (4)$$

where  $q \neq 1$  is a sensitivity parameter [41] that controls how much the diversity measure emphasizes common versus rare industries. When  $q$  is small,  ${}^q D(a, Z)$  gives as much importance to rare industries as common ones [41]; thus,  ${}^0 D(a, Z)$  is a measure of industry “richness” (the effective number of industries). By contrast, when  $q$  is large, rare industries are de-emphasized and  ${}^q D(a, Z)$  includes information about the evenness of industries.

## 6 Model plausibility results

We examine the firm-pair similarity matrices produced by the Boolean, TF-IDF, and PVDMM models, and compare them for plausibility by comparison with the simple SIC model’s similarity matrix (visible in Fig 4). Next, we test the plausibility of each model by seeing if they put similar firms in clusters, and if they give especially high similarity to pairs of firms with the same SIC Codes.

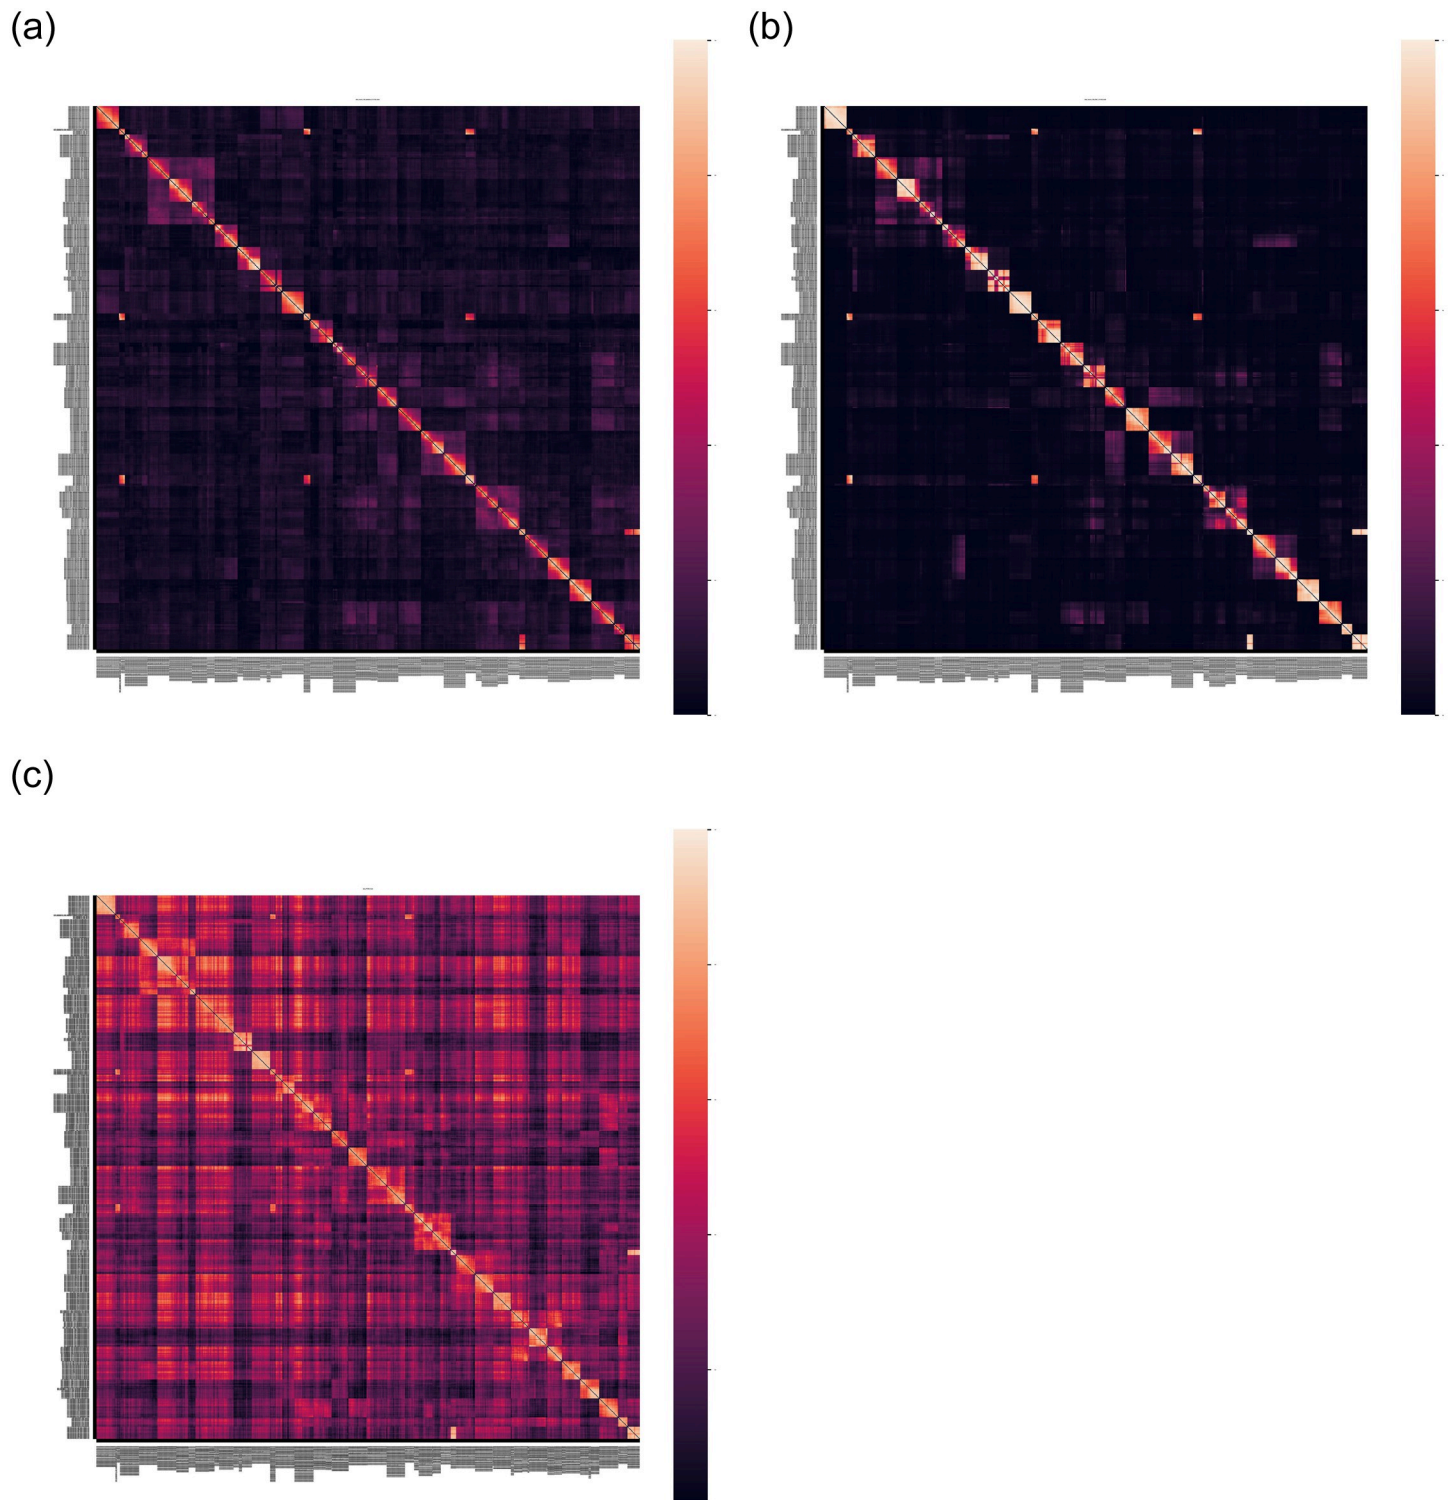
### 6.1 Firm similarity matrices

Fig 8 shows heatmaps of the similarity matrices of twenty years of product descriptions by the firms listed in the 1997 Dow Jones Industrial Average (DJIA), according to the text-based (Boolean, TF-IDF, PVDMM) models of firm level similarities. All of the heatmaps show high similarity among firms with the same SIC Codes (squares of high heat along the diagonal), faintly reflecting the SIC model’s white squares along the diagonal in Fig 4, and this pattern is quantitatively corroborated below with Industry Specificity measurements.

Furthermore, the firm similarity heatmaps shown in Fig 8 all show many differences in the similarity of different pairs of firms within the same SIC Code or higher-level SIC group. By contrast, the text-free SIC Code model depicts every pair of firms in each SIC group with the exact same level of similarity (Fig 4). The heatmaps contain a wealth of information about the different degrees of similarity attributed by each text-based model to each pair of individual firms.

### 6.2 Micro-analysis of clusters

To evaluate the text-based models and their consequent document embeddings, we focus on twenty five well-known firms: Boeing, Caterpillar, Cisco, Chevron, Disney, DuPont, General Electric, Home Depot, IBM, Intel, Johnson & Johnson, Coca-Cola, McDonalds, 3M, Merck, Microsoft, Nike, Pfizer, Procter & Gamble, AT&T, United Technologies, Visa, Verizon, Wal-



**Fig 8.** Heatmaps of firm similarity matrices of the products of the firms in the 2018 DJIA, according to the BOOLEAN (top left), TF-IDF (top right), and PVDM (bottom left) models. Compare with the SIC model similarity matrix shown in Fig 4.

<https://doi.org/10.1371/journal.pone.0264330.g008>



Mart, Exxon Mobil. All of these firms have been on the Dow Jones Industrial Average for much of this century. We evaluate the initial plausibility of each text-based model by seeing how well the proximity of each model's vectors for (documents about the products of) the twenty-five firms.

We gauge the proximity of embedded documents in high-dimensional vector spaces using 2-dimensional t-SNE projections [42] of the embedded documents. The t-SNE projections in Fig 9 depict the location in product space of documents about our twenty-five reference firms, according to the Boolean (left), TF-IDF (middle), and PV-DM (right) models. Each point indicates the location of an individual document in a given year, and the dots are color-coded to by firm. CIKs label each document cluster.

A striking feature of Fig 9 is that each cluster in the t-SNE projections contains documents from exactly one firm (one color); and documents about the products of different firms are in different clusters. This clear pattern is exactly what common-sense would expect from an embedding that reflects the real similarities among the products, and each text-based model produces this pattern.

Common-sense also suggests that some firms make quite different products. For example, the products of the following firms are all relatively distinctive and different from each other: McDonalds (63908), Coca-Cola (21344), Nike (320187), Disney (1001039), Visa (1403161). Note that Fig 9 shows that the Boolean, TF-IDF, and PV-DM models all place McDonalds, Coca-Cola, Nike, Disney, and Visa in relatively isolated locations in their respective product spaces.

Common sense also suggests that a plausible document embedding would put firms producing very similar products in nearby or overlapping clusters. For example, one would expect to see groups of nearby clusters (individual firms) containing the following groups of nearby DJIA firms (and their CIK number):

- Johnson & Johnson (200406), Merck (310658) and Pfiser (78003)
- Boeing (12927) and United Technologies (101829)
- AT&T (732717) and Verizon (732712)
- Home Depot (354950) and Walmart (104169)
- Exxon Mobile (34088) and Chevron (93410)
- IBM (51143), Intel (50863), Cisco (858877) and Microsoft (789019)

Inspection of the t-SNEs in Fig 9 confirms that the Boolean, TF-IDF, and PV-DM models all pass this additional test of common-sense. For ease of identification, the first three of these groups of firms are circled red, yellow and green.

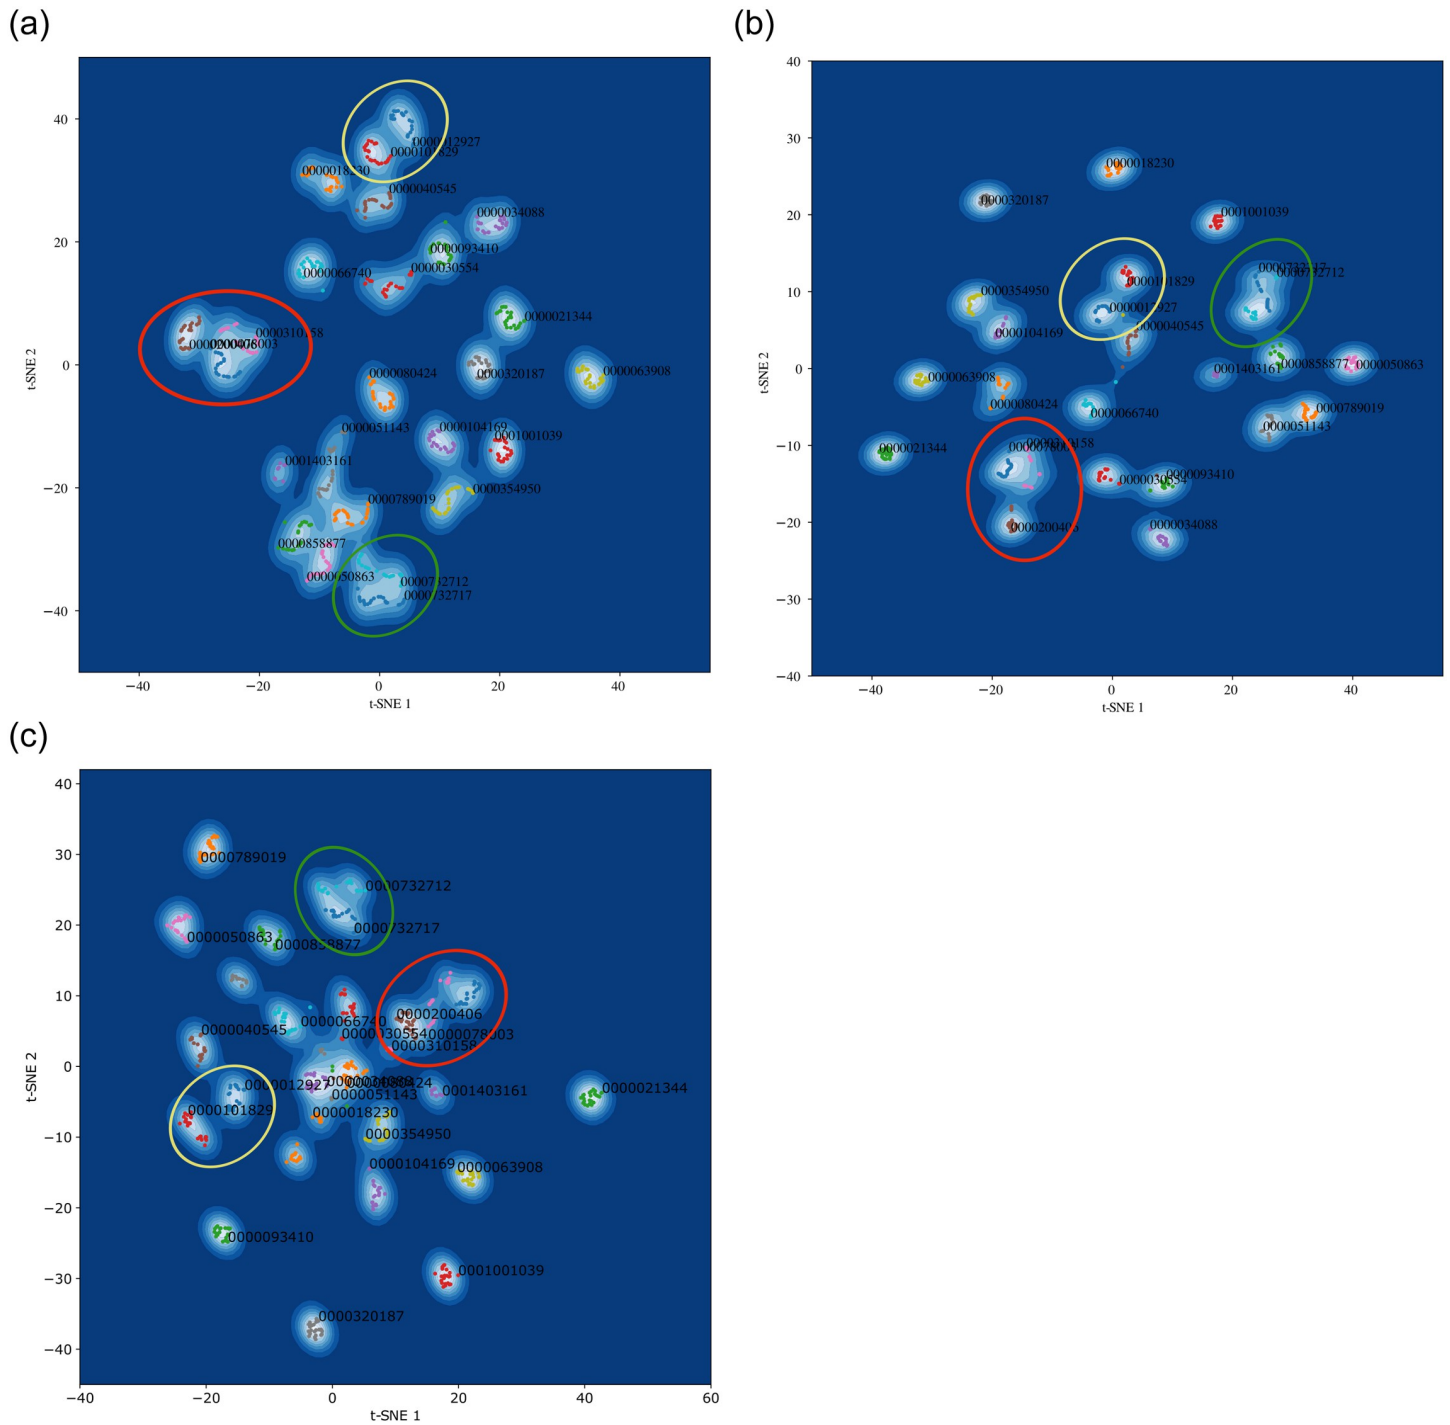
This micro-analysis of the details of the embeddings of firms in the DJIA adds weight to the general plausibility of all three document embeddings studied here. The Boolean, TF-IDF, and PV-MD models all demonstrate a significant degree of common-sense realism and plausibility.

### 6.3 SIC Industry Specificity results

The *SIC Industry Specificity* of each model (Boolean, TF-IDF, and PV-MD) is indicated in Fig 10. A vertical dotted line in the Figure shows the Industry Specificity of a Flat control model in which all firms are identical to one another. A model passes the SIC Industry Specificity test for plausibility if and only if it's Industry Specificity is significantly higher than the Flat control.

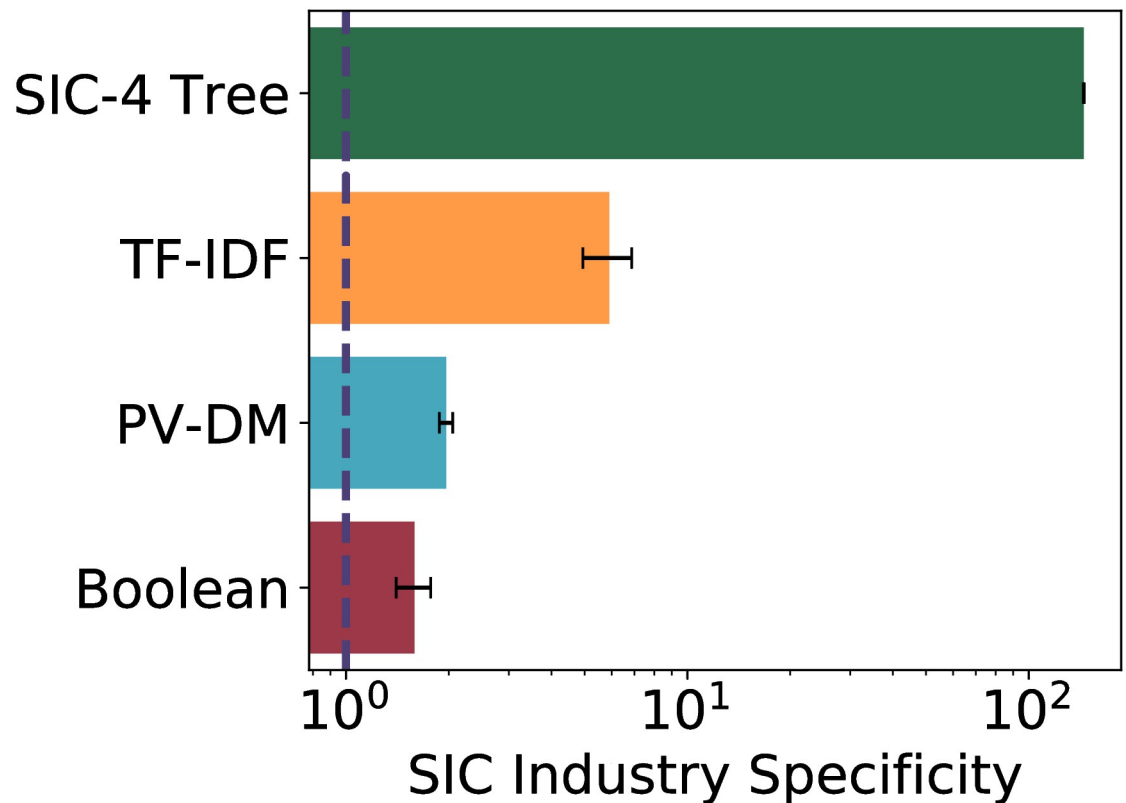
Fig 10 shows that the Specificity of the Boolean, TF-IDF, and PV-MD models all is significantly higher than the Flat control. The three text-based models all construe firms from the





**Fig 9.** t-SNE of documents describing the products of twenty-five firms from the Dow Jones Industrial Average, embedded in the product spaces from models Boolean (left), TF-IDF (middle), and PV-DM (right). Points indicate the location of a given document, colored by firm; document clusters are labeled by dominant CIK. Firms that produce pharmaceuticals are circled in red, aerospace companies are circled in yellow and telecommunication companies are circled in green.

<https://doi.org/10.1371/journal.pone.0264330.g009>



**Fig 10. Industry Specificity with respect to SIC of text-based models (Boolean, TFIDF, and PVDM) and the text-free SIC control.** The dotted line shows the null hypothesis of a perfectly flat similarity matrix.

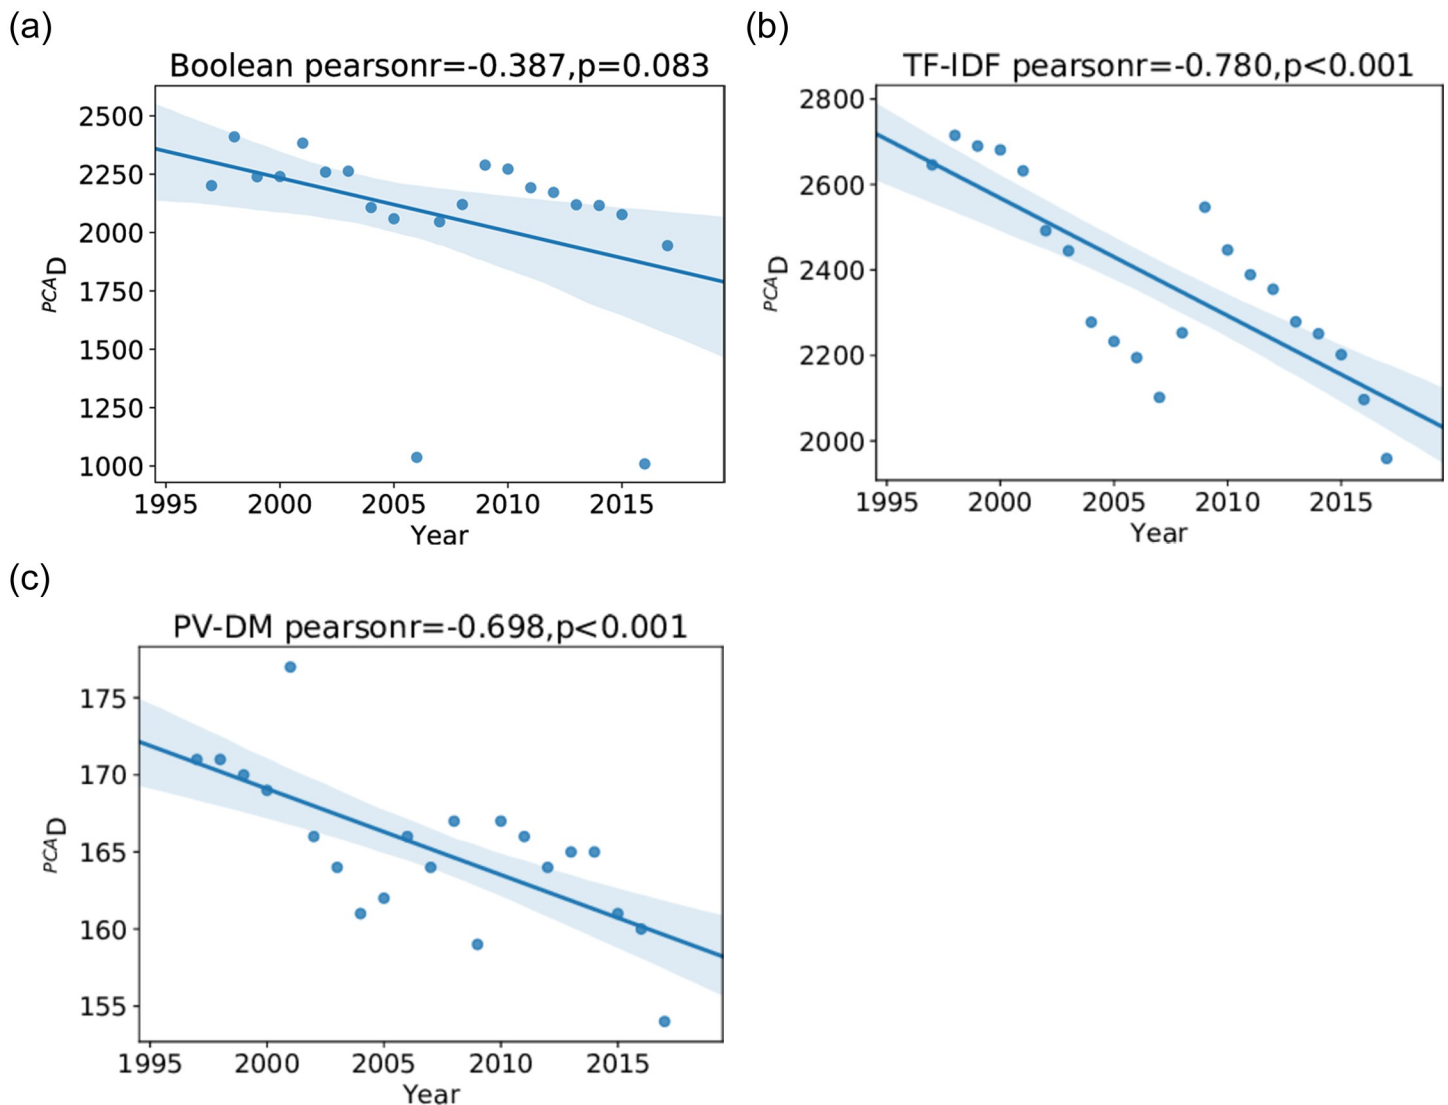
<https://doi.org/10.1371/journal.pone.0264330.g010>

same SIC class as more similar on average than firms from different classes. All of the text-based embeddings have higher firm similarities on average for pairs of firms with the same SIC Code. Thus, if the SIC in fact reflects some genuine structure among all the industries, the text-based vector embeddings reflect a similar structure.

Although a plausible model must have significantly more Industry Specificity than the Flat control, higher industry specificity is not always better. A model's Industry Specificity reflects the degree to which its similarity matrix corresponds to some trusted reference classification (here, SIC Codes). But there is no guarantee that the reference classification captures all the relevant observable details about the actual similarities among firms. A high Industry Specificity could reflect a model's high correspondence to a trusted reference's crude model of firm similarity. Furthermore, a model's level of Industry Specificity is roughly correlated with the off-diagonal heat in its firm similarity matrix, shown in Fig 8. But there is no reason to think that better models have lower off-diagonal heat, for there certainly is some similarity of some firms that are far apart in the SIC classification tree. So, any plausible model should have an SIC Industry Specificity that is significantly higher than the Flat control, but Industry Specificity should not be viewed as a model success metric.

## 7 Product diversity trends

We measure the diversity of the products of large US firms from the past twenty years (1997 to 2017), and look for trends over these years. We have already measured the annual diversity of products of large US firms simply as the number of four-digit SIC Codes that are instantiated



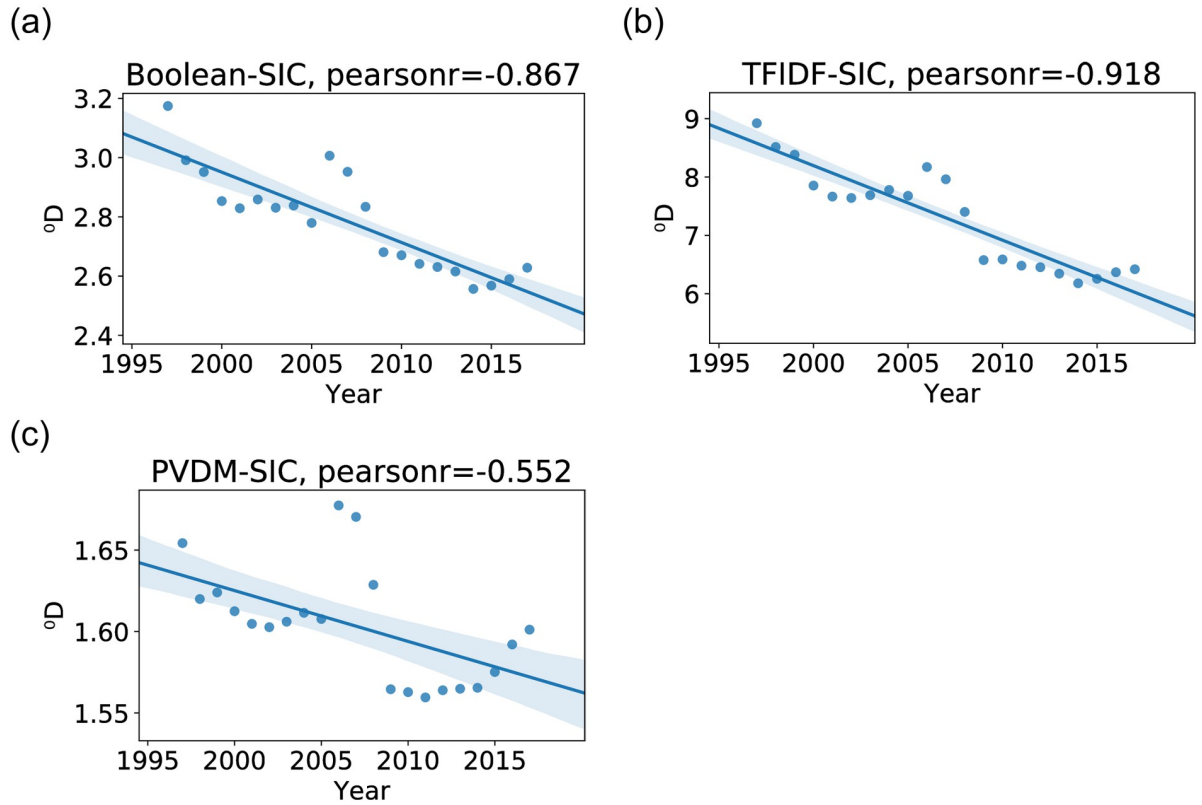
**Fig 11.** Annual number of dimensions required to account for 90% of variance of document vectors when embedded in the Boolean (left), TF-IDF (middle), and PV-DM (right) models. For each model we show a scatter plot of annual diversities and a linear regression fit with its 90% confidence interval.

<https://doi.org/10.1371/journal.pone.0264330.g011>

each year (Fig 7 left) and as the Shannon entropy of the distribution of SIC Code counts each year (Fig 7 right). We now focus on diversity measurements that depend on annual individual firm-level similarity matrices, produced by embedding in a semantic vector space a description of each firm's products in a given year. We compare the trends in product diversity disclosed by embedding firms in the Boolean, TF-IDF, and PV-DM vector spaces.

One measure of diversity,  $^{PCA}D$ , is simply the number of principle components required to account for 90% of the variance in the spread of firms embedded each year. Fig 11 shows  $^{PCA}D$  results for all three models, illustrating a consensus drop in diversity over the period of analysis. The Boolean and TF-IDF models show almost a 30% drop in  $^{PCA}D$  over twenty years, while the drop shown by the PV-DM model is only 8%.

In order to understand the effect on diversity trends of varying the degree of sensitivity to rare species, annual  $^qD$  values are calculated for  $q \in \{0, 2, 5\}$ . These diversity values reflect not just the abundances of different SIC classes but also how similar the classes are to each other.



**Fig 12.** Annual economic diversity of order zero,  ${}^0D$ , according to models Boolean (left), TF-IDF (middle), and PV-DM (right). Each model includes a scatter plot of diversities and a linear regression fit along with its 90% confidence interval.

<https://doi.org/10.1371/journal.pone.0264330.g012>

The scatterplots of annual diversity values and linear regression fits for  $q = 0$  and the three models of interest are shown for on Fig 12 while the Pearson correlation coefficients for all the tested sensitivities are shown on Table 1.

Table 1 and Fig 12 show that for  $q = 0$ , all three models show statistically significant decreasing trends in diversity. This means that all three models are in agreement that the richness of products is decreasing over the years. In other words, *the trend of dropping product richness* in the (descriptions of) products of large US firms is a consensus conclusion of diversity measurements with  $q = 0$  of Boolean, TF-IDF and PV-DM models. Similarly, all three text-based models show statistically significant decreasing trends in  ${}^2D$ .  ${}^2D$  starts to pay less emphasis to rare species and is equivalent to a commonly used diversity measure in ecology known as Rao’s quadratic entropy [41]. Finally, further increasing  $q$  to 5 continues the pattern of Boolean, TF-IDF and PV-DM showing statistically significant decreasing trends in diversity. The upshot of the diversity correlation coefficients is that all models show statistically significant patterns of dropping diversity across different sensitivity values. As with Shannon entropy, to

**Table 1. Correlation coefficients of diversity  ${}^qD$  with year and significance levels (\*\*: p-value  $\leq 0.05$ , \*\*\*: p-value  $\leq 0.01$ ).**

	${}^0D$	${}^2D$	${}^5D$
Boolean	-0.867***	-0.856***	-0.820***
TF-IDF	-0.918***	-0.894***	-0.808***
PV-DM	-0.552***	-0.543**	-0.531**

<https://doi.org/10.1371/journal.pone.0264330.t001>

take away the effect of the decreasing number of SIC Codes on  $^qD$ , the metric can be normalized as described in S2 Appendix in [S1 File](#). It can be seen that the normalized  $^0D$  also shows a decreasing trend in the Boolean and TF-IDF models and no significant trend with PV-DM.

## 8 Conclusions and discussion

This paper presents a wealth of evidence for a significant drop in the diversity of the products produced by large US firms in this century. This downward trend is evident whether diversity is measured in crude or sophisticated ways, and whether the information about the products of individual firms is coarse- or fine-grained. This trend can be seen using a Boolean word vector model, the current industry standard in product-focused firm embeddings due to Hoberg and Phillips [17], and the trend can be seen using more sophisticated TF-IDF and PV-DM models. The trend is even evident in a simple model based merely on a firm's four-digit SIC Code. The magnitude of the drop in diversity ranges from 6% to 30% depending on the method by which diversity is measured, and all the diversity measurements show some scatter year-by-year diversity, but the overall twenty-year trend of dropping product diversity is a very robust result.

Our product diversity results focus on large US firms, because our models are trained on documents that are filed only by large US firms. Since large US firms are an unrepresentative sample of all of the firms that contribute to the economy, so whether the dropping product diversity trend also holds for smaller firms and firms outside the US remains an open question. Even if we restrict our attention to large US firms, it also remains an open question how to explain the dropping diversity trend. We noted earlier an overall drop in number of firms over the same years (recall [Fig 3](#)), and this drop in the number of firms might be thought to explain the drop in diversity of products. Further, it is known that since the 1990s market concentration has been occurring as fewer firms take up more market share in their industries [43, 44]. However, we still observed the diversity drop when we measured diversity using normalized abundance vectors, so the drop in the number of firms is unlikely by itself to explain the observed trend in dropping product diversity.

A second, quite different hypothesis is document homogenization, which proposes that the decreasing diversity of the descriptions of firms' products is due merely to an increasing professionalization and standardization of the text in 10-K documents. This hypothesis suggests that models trained on textual documents provide evidence for a drop in diversity of the *descriptions* of products of large US firms, but *not* evidence for a decrease in diversity of the products themselves. The document homogenization hypothesis does not explain why the model of firm similarity based solely on a firm's SIC Code shows a similar drop in product diversity (recall [Fig 7](#)); nor does it explain the roughly 50% increase in average number of word tokens and word types in each document in the training corpus (recall [Fig 6](#)). So the dropping diversity seen using text-based models is unlikely to be due specifically to document homogenization.

A number of further hypotheses could explain the dropping product diversity trend. One is the hypothesis that products have shrunk in diversity because consumer demand for products has narrowed. Another hypothesis is that the growing diffusion of information technology into more and more products is making products overall more alike. A third hypothesis is that the drop in product diversity is due to the rise of outsourcing by large US firms, and a consequent *rise* in the diversity of products produced *outside* the United States. A fourth hypothesis would connect the drop in diversity of the products of *large* firms with a rise in the diversity of products produced by *small* firms. We have no specific evidence for or against any of these hypotheses, but all of them have empirically testable consequences.

However, gathering accurate and complete data about the products of firms of most sizes in most countries remains a huge hurdle.

One final hypothesis worth considering is that the trend of falling product diversity is explained by an increasing diversity of products *within* large US firms. On this hypothesis, the total diversity of products in the marketplace may be stable or growing, because individual large US firms on average have been producing an increasingly diverse array of products. The diversity of products produced by some individual firms has been studied, and some have grown more diverse over time. When we measure the diversity of the products produced by large US firms, the products of each firm is embedded as a point in a high-dimensional product space, and we measure the diversity of those points in product space. So, those measurements reflect the diversity *between* the products produced by different firms. Since all the products of an individual firm in a given year are embedded into a point in product space, the diversity measurements are blind to the diversity of products *within* each firm in that year. (Of course, the diversity measures do reflect changes over the years in the diversity of the product offerings of each individual firm. Recall the clusters of points with the same color in Fig 9.) Nevertheless, Fig 6 shows that there is a significant rise in the number of word tokens and word types in 10-K *descriptions* of the products of large US firms. This does provide some corroboration for the hypothesis that the products of individual large US firms has grown more diverse during the past twenty years.

## Supporting information

**S1 File.**  
(PDF)

## Acknowledgments

We thank Francois Lafond, Kieran Marray, Norman Packard, Doyne Farmer and Sergei Maslov for suggestions on the manuscript, and we thank Amir Amel-Zadeh for help with document pre-processing. This project was accomplished with the support of project “Predicting Technological Progress by Embedding Input-Output Networks in Semantic Technological Spaces” (PR0040), in response to IARPA-BAA-17-01.

## Author Contributions

**Conceptualization:** Ananthan Nambiar, Tobias Rubel, Mark Bedau.

**Data curation:** Tobias Rubel.

**Formal analysis:** Ananthan Nambiar, Tobias Rubel, Jon deVries.

**Funding acquisition:** Mark Bedau.

**Investigation:** Ananthan Nambiar, Tobias Rubel, James McCaull, Jon deVries.

**Methodology:** Ananthan Nambiar, Tobias Rubel, Mark Bedau.

**Project administration:** Tobias Rubel, Mark Bedau.

**Resources:** Mark Bedau.

**Software:** Ananthan Nambiar, James McCaull, Jon deVries.

**Supervision:** Mark Bedau.

**Validation:** Ananthan Nambiar, Tobias Rubel.



**Visualization:** Ananthan Nambiar, Tobias Rubel, James McCaull, Jon deVries.

**Writing – original draft:** Ananthan Nambiar, Tobias Rubel, Jon deVries, Mark Bedau.

**Writing – review & editing:** Ananthan Nambiar, Tobias Rubel, Jon deVries, Mark Bedau.

## References

1. Wagner JE, Deller SC. A measure of economic diversity: An input-output approach. *Staff Paper*. 1993;93.
2. Siegel P, Johnson T, Alwang J. Regional Economic Diversity and Diversification. *Growth and Change*. 1995; 26:261–284. <https://doi.org/10.1111/j.1468-2257.1995.tb00171.x>
3. Wagner JE. Regional economic diversity: action, concept, or state of confusion. *Journal of Regional analysis and policy*. 2000; 30(1100-2016-90005).
4. Stirling A. A general framework for analysing diversity in science, technology and society. *Journal of the Royal Society Interface*. 2007; 4(15):707–719. <https://doi.org/10.1098/rsif.2007.0213> PMID: 17327202
5. van Dam A. Diversity and its decomposition into variety, balance and disparity. *Royal Society open science*. 2019; 6(7):190452. <https://doi.org/10.1098/rsos.190452> PMID: 31417744
6. Dissart JC. Regional economic diversity and regional economic stability: research results and agenda. *International Regional Science Review*. 2003; 26(4):423–446. <https://doi.org/10.1177/0160017603259083>
7. Freire C. Economic Diversification: Explaining the pattern of diversification in the global economy and its implications for fostering diversification in poorer countries. *United Nations DESE Working Paper No 150*. 2017.
8. Tallman S, Li J. Effects of international diversity and product diversity on the performance of multinational firms. *Academy of Management Journal*. 1996; 39(1):179–196. <https://doi.org/10.5465/256635>
9. Anderson SP, Renault R. Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model. *The RAND Journal of Economics*. 1999; p. 719–735. <https://doi.org/10.2307/2556072>
10. Dixit AK, Stiglitz JE. Monopolistic competition and optimum product diversity. *The American economic review*. 1977; 67(3):297–308.
11. Yang X, Heijdra BJ. Monopolistic competition and optimum product diversity: Comment. *The American Economic Review*. 1993; 83(1):295–301.
12. d'Aspremont C, Ferreira RDS, Gérard-Varet LA. On the Dixit-Stiglitz model of monopolistic competition. *The American Economic Review*. 1996; p. 623–629.
13. Bu Y, Li M, Gu W, Huang Wb. Topic diversity: A discipline scheme-free diversity measurement for journals. *Journal of the Association for Information Science and Technology*. 2021; 72(5):523–539. <https://doi.org/10.1002/asi.24433>
14. Elejalde E, Ferres L, Herder E, Bollen J. Quantifying the ecological diversity and health of online news. *Journal of Computational Science*. 2018; 27:218–226. <https://doi.org/10.1016/j.jocs.2018.06.003>
15. Mealy P, Farmer JD, Teyrelboym A. Interpreting economic complexity. *Science Advances*. 2019;(5): eaau1705. <https://doi.org/10.1126/sciadv.aau1705> PMID: 30662945
16. Hoberg G, Phillips G. Product market synergies and competition in mergers and acquisitions: A text-based analysis. *The Review of Financial Studies*. 2010; 23(10):3773–3811. <https://doi.org/10.1093/rfs/hhq053>
17. Hoberg G, Phillips G. Text-based network industries and endogenous product differentiation. *Journal of Political Economy*. 2016; 124(5):1423–1465. <https://doi.org/10.1086/688176>
18. Hoberg G, Phillips G, Prabhala N. Product market threats, payouts, and financial flexibility. *The Journal of Finance*. 2014; 69(1):293–324. <https://doi.org/10.1111/jofi.12050>
19. Hoberg G, Maksimovic V. Redefining financial constraints: A text-based analysis. *The Review of Financial Studies*. 2015; 28(5):1312–1352. <https://doi.org/10.1093/rfs/hhu089>
20. Hoberg G, Lewis C. Do fraudulent firms produce abnormal disclosure? *Journal of Corporate Finance*. 2017; 43:58–85. <https://doi.org/10.1016/j.jcorpfin.2016.12.007>
21. Bellstam G, Bhagat S, Cookson JA. A text-based analysis of corporate innovation. *Management Science*. 2020.
22. Hanley KW, Hoberg G. The information content of IPO prospectuses. *The Review of Financial Studies*. 2010; 23(7):2821–2864. <https://doi.org/10.1093/rfs/hhq024>
23. Agarwal S, Gupta S, Israelsen RD. Public and private information: Firm disclosure, sec letters, and the jobs act. *Georgetown McDonough School of Business Research Paper*. 2017;(2891089):17–4.

24. Lowry M, Michaely R, Volkova E. Information Revealed through the Regulatory Process: Interactions between the SEC and Companies ahead of Their IPO. *The Review of Financial Studies*. 2020; 33(12):5510–5554. <https://doi.org/10.1093/rfs/hhaa007>
25. Popadak J. A corporate culture channel: How increased shareholder governance reduces firm value. Available at SSRN. 2013;2345384.
26. Israelsen RD. Does common analyst coverage explain excess comovement? *Journal of Financial and Quantitative Analysis*. 2016; p. 1193–1229. <https://doi.org/10.1017/S002210901600051X>
27. Goldsmith-Pinkham P, Hirtle B, Lucca DO. Parsing the content of bank supervision. 2016.
28. Bergeaud A, Potiron Y, Raimbault J. Classifying patents based on their semantic content. *PLoS ONE*. 2017;(12):e0176310. <https://doi.org/10.1371/journal.pone.0176310> PMID: 28445550
29. Bedau M, Gigliotti N, Janssen T, Kosik A, Nambiar A, Packard N. Open-ended technological innovation. *Artificial Life*. 2019;(25):33–49. [https://doi.org/10.1162/artl\\_a\\_00279](https://doi.org/10.1162/artl_a_00279) PMID: 30933632
30. Tshitoyan V, Dagdelen J, Weston L, Dunn A, Rong Z, Kanonova O, et al. Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*. 2019;(571):95–98. <https://doi.org/10.1038/s41586-019-1335-8> PMID: 31270483
31. Rives A, Meier J, Sercu T, Goyal S, Lin Z, Liu J, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*. 2021; 118(15). <https://doi.org/10.1073/pnas.2016239118> PMID: 33876751
32. Nambiar A, Heflin M, Liu S, Maslov S, Hopkins M, Ritz A. Transforming the Language of Life: Transformer Neural Networks for Protein Prediction Tasks. In: *Proceedings of the 11th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics. BCB'20*. New York, NY, USA: Association for Computing Machinery; 2020. Available from: <https://doi.org/10.1145/3388440.3412467>.
33. Elnaggar A, Heinzinger M, Dallago C, Rihawi G, Wang Y, Jones L, et al. ProtTrans: Towards Cracking the Language of Life's Code Through Self-Supervised Deep Learning and High Performance Computing. *CoRR*. 2020;abs/2007.06225.
34. Standard Industrial Classification (SIC);. Available from: [https://www.osha.gov/pls/imis/sic\\_manual.html](https://www.osha.gov/pls/imis/sic_manual.html).
35. Fixler D, Morgan E, Murphy J, Talan D, Bugg P. North American Industry Classification System (NAICS); 2017. Available from: [https://www.census.gov/eos/www/naics/2017NAICS/2017\\_NAICS\\_Manual.pdf](https://www.census.gov/eos/www/naics/2017NAICS/2017_NAICS_Manual.pdf).
36. SEC.gov j Form 10-K; 2009. Available from: <https://www.sec.gov/fast-answers/answers-form10khtm.html>.
37. Loughran T, McDonald B. Textual Analysis in Accounting and Finance: A Survey. *SSRN, Journal of Accounting Research*. 2016; 54:4:1187–1230. <https://doi.org/10.1111/1475-679X.12123>
38. Robertson S. Understanding inverse document frequency: on theoretical arguments for IDF. *Journal of Documentation*. 2004; 60:503–520. <https://doi.org/10.1108/00220410410560582>
39. Le Q, Mikolov T. Distributed Representations of Sentences and Documents. In: Xing EP, Jebara T, editors. *Proceedings of the 31st International Conference on Machine Learning*. vol. 32 of *Proceedings of Machine Learning Research*. Beijing, China: PMLR; 2014. p. 1188–1196. Available from: <http://proceedings.mlr.press/v32/le14.html>.
40. Řehůřek R, Sojka P. Software Framework for Topic Modelling with Large Corpora. In: *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. Valletta, Malta: ELRA; 2010. p. 45–50. Available from: <http://is.muni.cz/publication/884893/en>.
41. Leinster T, Cobbold CA. Measuring diversity: the importance of species similarity. *Ecology*. 2012; 93(3):477–489. <https://doi.org/10.1890/10-2402.1> PMID: 22624203
42. van der Maaten L, Hinton G. Visualizing data using t-SNE; 2008.
43. Grullon G, Larkin Y, Michaely R. Are US industries becoming more concentrated? *Review of Finance*. 2019; 23(4):697–743. <https://doi.org/10.1093/rof/rfz007>
44. Bessen J. Information technology and industry concentration. 2017.